CFO: Calibration-Free Odds Bayesian Designs for Dose Finding in Clinical Trials

Jialu Fang, MS[1]; Ninghao Zhang, MS[1]; Wenliang Wang, MPhil[1]; and Guosheng Yin, PhD[1]

DOI https://doi.org/10.1200/CCI-24-00184

ABSTRACT

PURPOSE Calibration-free odds type (CFO-type) designs have been demonstrated to be robust, model-free, and practically useful, which have become the state-of-the-art approach for dose finding. However, a key challenge for implementing such designs is a lack of accessible tools. We develop a user-friendly *R* package and *Shiny* web-based software to facilitate easy implementation of CFO-type designs. Moreover, we incorporate randomization into the CFO framework.

METHODS We created the *R* package CFO and leveraged *R Shiny* to build an interactive web application, CFO suite, for implementing CFO-type designs. We introduce the randomized CFO (rCFO) design by integrating the exploration-exploitation mechanism into the CFO framework.

RESULTS The CFO package and CFO suite encompass various variants tailored to different clinical settings. Beyond the fundamental CFO design, these include the two-dimensional CFO (2dCFO) for drug-combination trials, accumulative CFO (aCFO) for accommodating all dose information, rCFO for integrating exploration-exploitation via randomization, time-to-event CFO (TITE-CFO), and fractional CFO (fCFO) for addressing late-onset toxicity. Using all information and addressing delayed toxicity outcomes, TITE-aCFO and fractional-aCFO are also included. The package provides functions for determining the subsequent cohort dose, selecting the maximum tolerated dose, and conducting simulations to evaluate performance, with results presented through textual and graphical outputs.

CONCLUSION The CFO package and CFO suite provide comprehensive and flexible tools for implementing CFO-type designs in phase I clinical trials. This work is highly significant as it integrates all existing CFO-type designs to facilitate novel trial designs with enhanced performance. In addition, this promotes the spread of statistical methods using a user-friendly *R* package and *Shiny* software. It strengthens collaborations between biostatisticians and clinicians, further enhancing trial performance in terms of efficiency and accuracy.

ACCOMPANYING CONTENT

Data Supplement

Accepted December 10, 2024
Published January 31, 2025

JCO Clin Cancer Inform
9:e2400184

INTRODUCTION

The primary objective of a phase I clinical trial in oncology is to determine the maximum tolerated dose (MTD), defined as the dose where the probability of dose-limiting toxicity (DLT) aligns with a predetermined toxicity rate.[1] Numerous dose-finding methodologies currently exist for MTD determination, categorized broadly into algorithm-based, model-based, and model-assisted approaches. The traditional 3 + 3 design,[2] while simple and transparent, sometimes struggles to accurately identify the MTD and tends to assign subtherapeutic doses.[3] Conversely, model-based approaches, like the continual reassessment method (CRM)[4] and the escalation with overdose control (EWOC) design,[5] provide more accurate dose adjustments but are sensitive to the parametric model assumptions, which, if violated, may impair performance. Model-assisted and certain algorithm-based designs, including the cumulative cohort design (CCD),[6] Bayesian optimal interval (BOIN) design,[7] uniformly most powerful Bayesian interval design,[8] and calibration-free odds (CFO) design,[9] seek to combine the simplicity of algorithm-based methods with the precision of model-based approaches. Specifically, the CFO design, which uses a Bayesian framework without explicit dose-

toxicity models, enhances robustness by leveraging data across multiple dose levels, unlike interval-based designs that focus solely on the current dose.

The CFO design,[9] a novel phase I trial methodology, has been demonstrated to be robust, model-free, and easy to use. With the emergence of the time-to-event (TITE) method and fractional method, the TITE-CFO[10] and fractional CFO (fCFO) designs are developed to accommodate delayed toxicity. In these CFO-type designs, only a subset (current dose and its two neighboring doses) of the complete dose information is used. To incorporate data from all dose levels, the accumulative CFO (aCFO) design[11] is proposed and further extended to handle late-onset toxicity as the TITE-aCFO and fractional-aCFO (f-aCFO) designs.[11] Recently, the two-dimensional CFO (2dCFO) design[12] has been developed to advance CFO for use in drug-combination trials. Extensive simulations indicate comparable, and at times superior, performance compared with competing methodologies.[9-11]

Existing phase I designs typically use greedy, deterministic strategies that exploit past information without exploring unknowns. We propose the randomized CFO (rCFO) design, which integrates the exploration-exploitation mechanism from reinforcement learning into the CFO framework. This introduces probabilistic dose adjustments, improving decision making while retaining the model-free and calibration-free nature of CFO and ensuring robustness and objectivity without requiring artificial parameter inputs.

Various related *R* packages have been developed for dose-finding methodologies. Examples of such packages include BOIN[13] and TITEgBOIN[14] for BOIN-type designs, bcrm[15] and dfcrm[16] for CRM-type designs, ewoc[17] for EWOC-type designs, and TEQR[18] for CCD-type designs. However, *R* packages for CFO-type designs are yet to be developed. In this article, we introduce a comprehensive, well-documented, and user-friendly *R*[19] package—the CFO package accompanied by its *Shiny* application.

The CFO package represents a comprehensive implementation of CFO-type designs for phase I trials, available from the Comprehensive *R* Archive Network.[20] It supports key functionalities such as determining the subsequent cohort's dose level, selecting the MTD for a single trial, and executing simulations to obtain the operating characteristics. CFO offers flexibility in selecting CFO-type designs tailored to specific trial needs, like incorporating all dose information, randomization, handling late-onset toxicity, and accommodating single or combination therapies. The functions for distinct tasks under various CFO-type designs are illustrated in Figure 1. CFO distinguishes itself by providing a user-friendly evaluation through both summary and graphical outputs via summary() and plot(). This feature provides users with a more intuitive understanding of the model's operational dynamics and outcomes, thereby facilitating broader utilization. In addition, the CFO suite, an interactive web application built using *R Shiny*, allows users to implement CFO-type designs directly via the webpage.[21]

The article is organized as follows: The Methods section introduces the innovative rCFO design and other existing CFO-type designs, with real application using *R*. The Simulations section explains conducting simulations in *R*. The Real Trial section illustrates the CFO suite with a trial example. The Summary section concludes with a discussion. Because of space constraints, usage instructions are kept concise, with more detailed guidance available in the Data Supplement.

METHODS

In toxicity monitoring, CFO-type designs aim to determine the MTD with a DLT risk closest to the target rate. Figure 2 presents a flowchart summarizing the sequence of steps. If stopping conditions are not met, odds ratios are calculated for designs without late-onset toxicity. For those with late-onset toxicity, pending data must be completed first. With the completed or inherently complete data, dose allocation is made using information embedded in a subset or at all dose levels. This iterative process continues until the sample size is exhausted, and the MTD is finally identified through isotonic regression.[22] Descriptions of the CFO package functions and their arguments are detailed in the Data Supplement.

The CFO and aCFO Designs for Single-Drug Trials

Suppose that a clinical trial examines K dose levels with DLT rates, $p_1 < \cdots < p_K$, and a target DLT rate of ϕ. After enrolling n cohorts, all relevant dose data are combined into the cumulative data set $D_n = (x_k, m_k)_{k=1}^{K}$, where the n-th cohort is treated at dose level C. There are J doses to the left and H doses to the right of C, satisfying $K = J + H + 1$. The DLT rates for all doses can also be denoted as $(p_{L_J}, \ldots, p_{L_1}, p_C, p_{R_1}, \ldots, p_{R_H})$, with $p_{L_J} = p_1$ and $p_{R_H} = p_K$.

The CFO design determines the most appropriate dose level for the next cohort by comparing the current dose level (C) with the dose one level to the right (R_1) and one level to the left (L_1). The odds of the true DLT rate p_k being greater than the target DLT rate ϕ is defined as

$$O_k = \frac{\Pr(p_k > \phi \mid x_k, m_k)}{\Pr(p_k \le \phi \mid x_k, m_k)}, \quad k = L_1, C, R_1.$$

The reciprocal $\bar{O}_k = 1/O_k$ represents the odds of $p_k \le \phi$.

To illustrate the decision for dose de-escalation, O_C measures evidence of excessive toxicity at the current dose level, with a

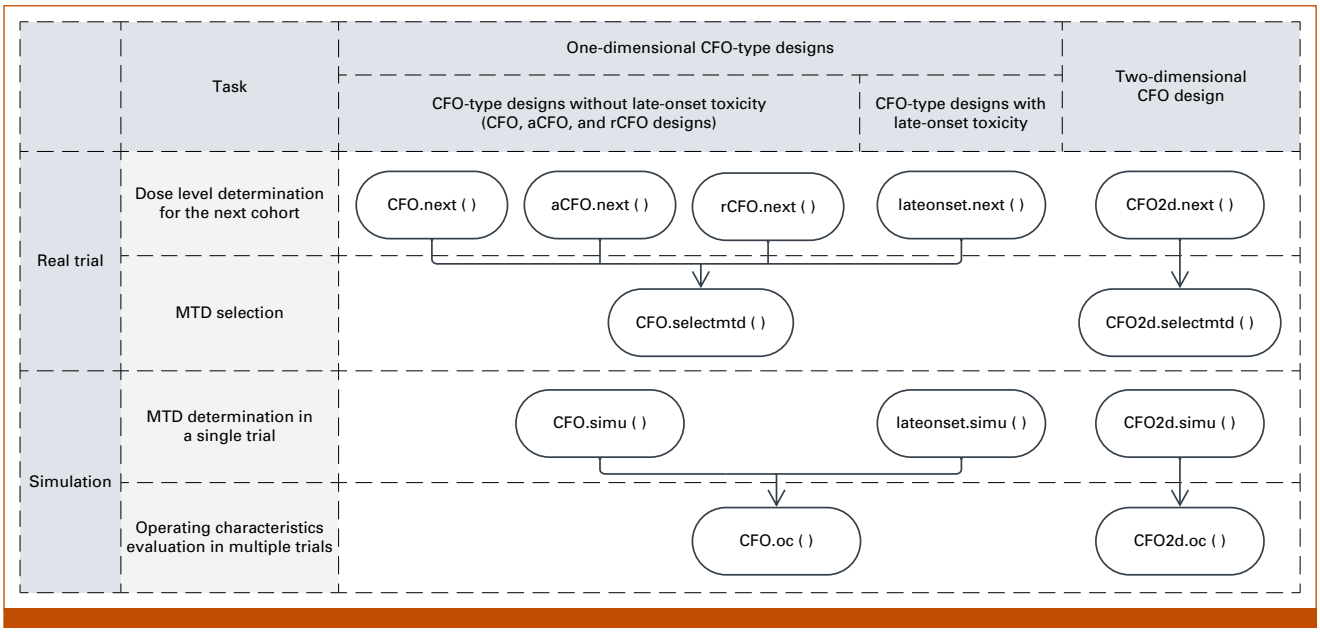

FIG 1. The utilization flowchart of user-visible functions in the CFO package. aCFO, accumulative CFO; CFO, calibration-free odds; MTD, maximum tolerated dose; rCFO, randomized CFO.

higher O_C indicating a preference for dose de-escalation. Conversely, a higher $\bar{O}_L$ reflects excessive tolerance at the left dose level, making de-escalation less favorable. The interplay between O_C and $\bar{O}_L$ results in the OR $O_C/\bar{O}_L$, indicating the strength of the inclination toward dose de-escalation. Similarly, $\bar{O}_C/O_R$ measures the strength toward escalation to the right dose level. The thresholds for the odds ratios, γ_L for de-escalation and γ_R for escalation, are predetermined by minimizing the risk of incorrect decisions. By integrating these decision processes, the dose level for the next cohort is determined following the decision rule for the CFO design specified in Table 1.

As the trial advances, it is crucial to consider the accumulative data at distant dose levels that could hold valuable information. The aCFO design incorporates this by accumulating information from all dose levels to the left (or right) of the current dose level.[11] Building upon the odds ratios $O_C/\bar{O}_L$ and $\bar{O}_C/O_R$, two aggregated OR statistics are formulated, encompassing comprehensive leftward and rightward information:

$$\mathrm{OR}_L = \frac{O_C}{\bar{O}_{L_1}} + \frac{O_C}{\bar{O}_{L_2}} + \ldots + \frac{O_C}{\bar{O}_{L_j}},$$

$$\mathrm{OR}_R = \frac{\bar{O}_C}{O_{R_1}} + \frac{\bar{O}_C}{O_{R_2}} + \ldots + \frac{\bar{O}_C}{O_{R_H}}.$$

This resembles the tug-of-war by summing strengths from the left versus those from the right. New thresholds are determined by summing their respective individual thresholds, and the decision rule for the aCFO design is outlined in Table 1.

In a hypothetical phase I trial with seven dose levels and a target DLT rate of 0.2, suppose that the current dose level is 3. To decide the next cohort's dose level using the CFO design, the function CFO.next() is executed:

```
R > decision <- CFO.next(target = 0.2, cys = c(0, 1, 0),
cns = c(3, 6, 0), currdose = 3, cutoff.eli = 0.95,
early.stop = 0.95)
R > summary(decision)
```

The process in the aCFO design resembles that of CFO but requires the use of aCFO.next().

Designs With Late-Onset Toxicities

Late-onset toxicity commonly arises in phase I dose-finding trials. The follow-up time for pending data contains rich information for refining dose selection. The TITE and fractional frameworks handle it by representing pending DLT data with decimal values between 0 (no DLT) and 1 (DLT occurrence). Designs accommodating late-onset toxicity, including TITE-CFO,[10] fCFO, TITE-aCFO, and f-aCFO,[11] have been proposed under these two frameworks. Specifically, the time-to-event framework assumes a uniform distribution for the time to DLT, whereas the fractional framework uses the Kaplan-Meier estimator without assumptions about the time-to-event data.

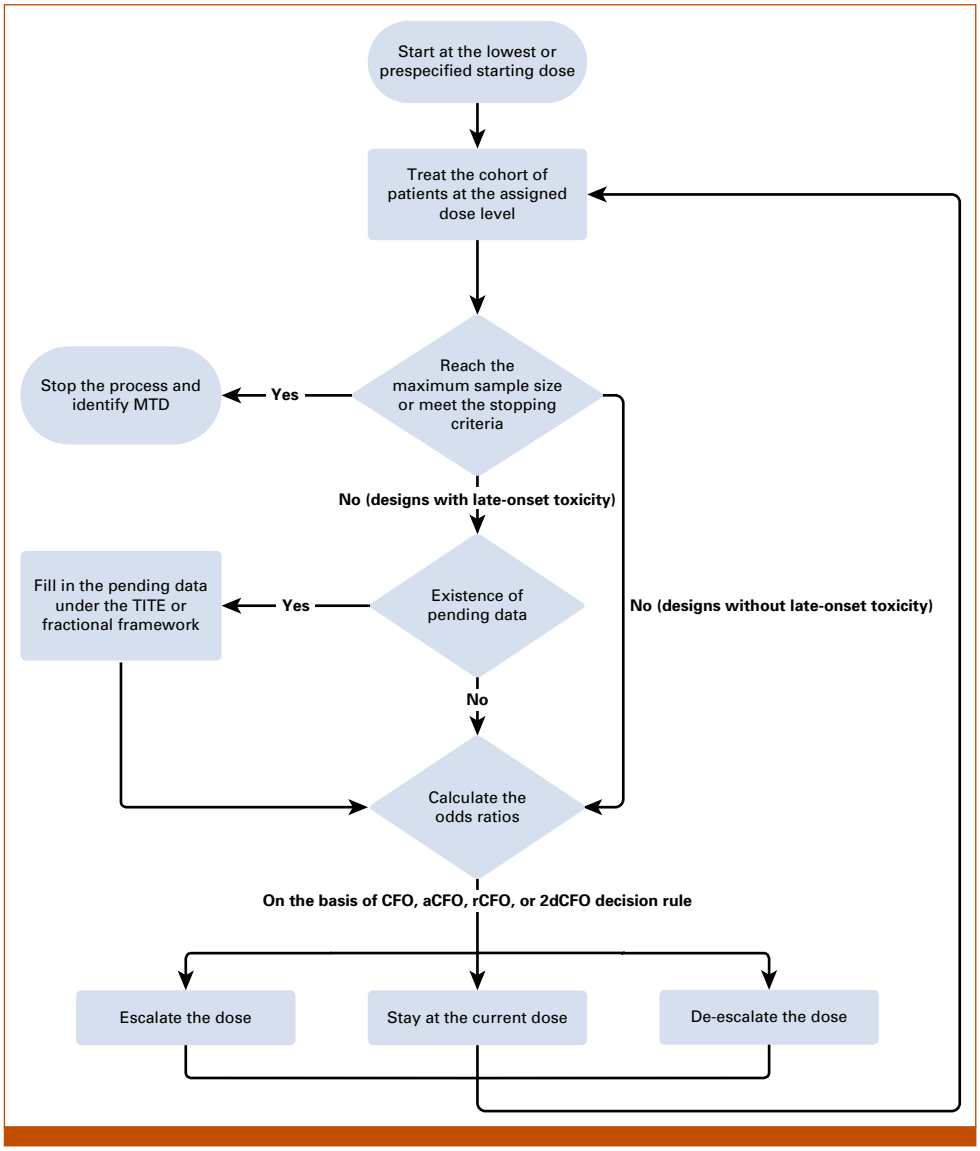

FIG 2. The flowchart of the Bayesian CFO-type design for phase I clinical trials. 2dCFO, two-dimensional CFO; aCFO, accumulative CFO; CFO, calibration-free odds; rCFO, randomized CFO; TITE, time-to-event.

The assessment window is denoted by τ, the follow-up time of a patient with the pending DLT outcome is u, and the time to DLT is represented by T. In the framework of the time-to-event weighting model, assuming a uniform distribution over the interval $[0, \tau]$, the TITE-CFO and TITE-aCFO designs address pending y by considering the expected outcome conditioned on the actual follow-up time. For a patient treated at dose level k, the imputed outcome is

$$\begin{aligned}\hat{y} &= \mathrm{E}(y|T>u) \\ &= \frac{\Pr(y=1)\Pr(T>u|y=1)}{\Pr(y=1)\Pr(T>u|y=1)+\Pr(y=0)\Pr(T>u|y=0)} \\ &= \frac{p_k(1-u/\tau)}{p_k(1-u/\tau)+(1-p_k)},\end{aligned}$$

where p_k represents the true DLT rate at dose level k.

In the fractional framework, the contribution of pending data is modeled using the Kaplan-Meier estimator. Both fCFO and f-aCFO designs estimate the conditional probability of toxicity occurrence in the remaining follow-up period, given that the toxicity event has not occurred by time u. This is formulated as

$$\hat{y} = \Pr(T<\tau \,|\, T>u) = \frac{\Pr(u<T<\tau)}{\Pr(T>u)} = \frac{\hat{S}(u)-\hat{S}(\tau)}{\hat{S}(u)},$$

where $\hat{S}(\cdot)$ denotes the Kaplan-Meier estimator for the survival function $S(\cdot)$.

When addressing late-onset toxicity, users can use lateonset.next() to assign the most appropriate dose to each cohort. Taking the f-aCFO design as an example, dose

TABLE 1. Decision Rules for the CFO, aCFO, and rCFO Designs in Searching for the Maximum Tolerated Dose

CFO Design

		$O_C / \bar{O}_L > \gamma_L$	
		Yes (de-escalation)	No (stay)
$\bar{O}_C / O_R > \gamma_R$	Yes (escalation)	Stay	Escalation
	No (stay)	Escalation	Stay

aCFO Design

		$\sum_{i=1}^{J} \frac{O_C}{\bar{O}_{L_i}} > \sum_{i=1}^{J} \gamma_{L_i}$	
		Yes (de-escalation)	No (stay)
$\sum_{i=1}^{H} \frac{\bar{O}_C}{O_{R_i}} > \sum_{i=1}^{H} \gamma_{R_i}$	Yes (escalation)	Stay	Escalation
	No (stay)	Escalation	Stay

rCFO Design

		$\pi_L = O_C / \bar{O}_L > \gamma_L$	
		Yes (de-escalation)	No (stay)
$\pi_R = \bar{O}_C / O_R > \gamma_R$	Yes (escalation)	Stay if $\pi_L = \pi_R$; otherwise, escalate with probability $\pi_R/(\pi_L + \pi_R)$ and de-escalate with probability $\pi_L/(\pi_L + \pi_R)$	Escalate with probability $\pi_R/(\pi_L + \pi_R)$ and stay with probability $\pi_L/(\pi_L + \pi_R)$
	No (stay)	De-escalate with probability $\pi_L/(\pi_L + \pi_R)$ and stay with probability $\pi_R/(\pi_L + \pi_R)$	Stay

Abbreviations: aCFO, accumulative CFO; CFO, calibration-free odds; rCFO, randomized CFO.

assignment for the newly cohort can be performed as follows:

```
R > decision <- lateonset.next(design = 'f - aCFO',
target = 0.2, ndose = ndose, currdose = 4,
assess.window = 3, enter.times = enter.times,
dlt.times = dlt.times, current.t = 9.41, doses = doses,
cutoff.eli = 0.95, early.stop = 0.95)
R > summary(decision)
```

The 2dCFO Design for Drug-Combination Trials

Combined drugs have become commonplace for cancer treatment. To enhance the robustness and precision of CFO in drug-combination trials, the 2dCFO approach is introduced.[12] Decision making within the two-dimensional toxicity probability space involves performing two one-dimensional CFO analyses along the horizontal and vertical axes. Let C be the current dose, with adjacent doses L, R, U, and D (left, right, up, and down). The 2dCFO design uses the same odds and threshold formulation as its 1dCFO counterpart. Subject to the constraints of partial ordering, decisions are made on the basis of the horizontal direction (L, C, R) and vertical direction (D, C, U). Both sequences have monotonically ascending DLT rates. Additional sequences (L, C, U) and (D, C, R) are considered when necessary. The CFO2d.next() function is used to determine the subsequent dose levels (see the Data Supplement for details).

The rCFO Design With Randomization for Single-Drug Trials

Existing phase I trial designs predominantly use greedy approaches, making deterministic decisions by exploiting data from treated patients without exploring unknowns. Traditional exploiting strategies typically focus solely on past information to maximize immediate outcomes, often leading to local optima. We incorporate the exploitation-exploration concept from reinforcement learning, balancing the use of existing information (exploitation) with the investigation of new possibilities (exploration) to gather additional insights. While exploration may not yield the best immediate results, it deepens understanding of the dose-response relationship and may lead to more favorable long-term outcomes.

Inspired by this idea, we propose the rCFO design, which has not been previously explored in the dose-finding literature. The rCFO design incorporates the exploitation-exploration mechanism into the CFO framework, enabling the utilization of existing data for decision making while allowing for the exploration of other potential dose levels with probabilities. The original CFO design determines dose movement by comparing two odds ratios, $\pi_L = O_C / \bar{O}_L$ and $\pi_R = \bar{O}_C / O_R$, against thresholds γ_L and γ_R, respectively. In the CFO

framework, dose adjustments are made deterministically. The rCFO design introduces a randomization scheme, akin to the idea in multiarmed bandit problems, allowing for probabilistic dose adjustments.

The symbols π_L and π_R denote the strength of the inclination toward dose de-escalation and escalation, respectively. For instance, when $\pi_L > \gamma_L$ and $\pi_R > \gamma_R$, indicating the presence of both de-escalation and escalation tendencies, the CFO design mitigates these opposing trends and opts to stay at the current dose level. This conflicting interplay does not result in random decisions but rather prompts a deterministic choice to retain the current dosage. In the rCFO design, the randomized mechanism is introduced, wherein the core process of deterministic decision making in CFO is transformed into probabilistic decision making. Specifically, the rCFO design normalizes odds ratios into probabilities, constructing randomization probabilities for dose escalation, de-escalation, and staying at the same dose. By transforming odds ratios into probabilities, the rCFO design can make randomized decisions on the basis of varying degrees of inclination. This strategic shift aims to strike a balance between dose escalation and de-escalation, rather than simply offsetting each other. As delineated in the decision rule for rCFO design in Table 1, this stochastic decision rule facilitates dose escalation, de-escalation, or staying at the same dose on the basis of calculated probabilities. rCFO.next() in the CFO package is used to determine the dose level for the next cohort (see the Data Supplement for details).

We conducted 5,000 simulations to compare the CFO and rCFO designs under five fixed scenarios from the study by Cheung and Chappell,[23] with a target DLT rate of 0.2. Detailed descriptions of the scenarios and performance metrics are provided in the Data Supplement. The operational characteristics of the rCFO design are evaluated against the CFO design with sample sizes of 36 and 99. Figure 3 summarizes the overall assessment of accuracy and safety using five performance metrics.

Simulation results indicate that when the sample size is 36, the rCFO and CFO designs show comparable efficiency and accuracy, with rCFO slightly underperforming in safety. However, when the sample size is 99 the rCFO design, with a higher value of MTD selection and allocation, marginally outperforms CFO in efficiency and accuracy. The rCFO design is more conservative than CFO in overdose selection and allocation, whereas the differences in the average DLT rate are negligible. With smaller sample sizes, the benefits of randomization are less pronounced because of higher variability, but with a larger sample size (eg, 99), randomization becomes more effective. Moreover, a notable observation is

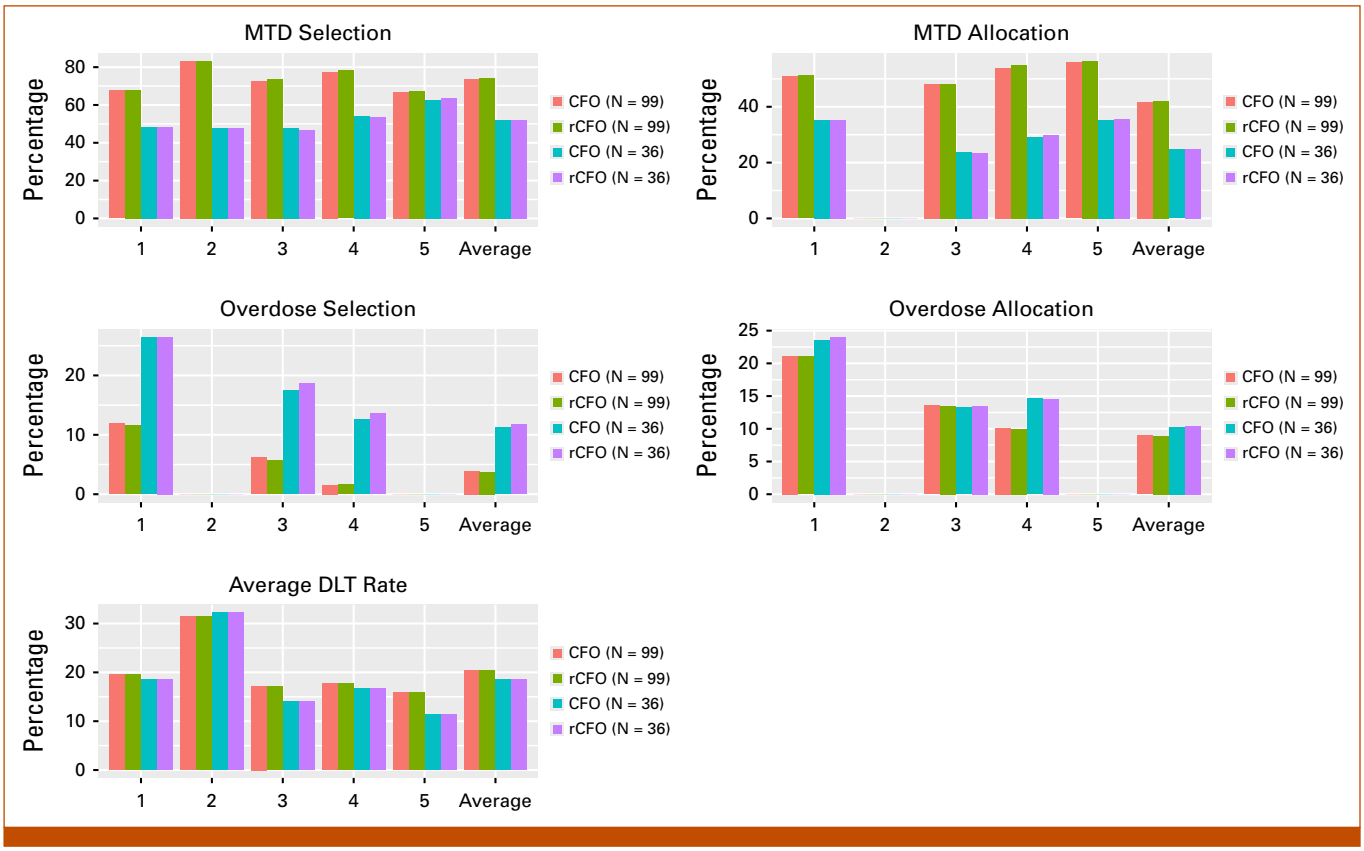

FIG 3. Simulation results of the CFO and rCFO designs with the target DLT rate of 0.2 and the sample size (N) of 36 and 99 under five fixed scenarios. The numbers 1–5 on the x-axis represent the five scenarios, and "Average" represents the metrics averaged across these scenarios. For MTD selection and allocation, a higher value is preferred. For overdose selection, overdose allocation, and the average DLT rate, a lower value is preferred. CFO, calibration-free odds; MTD, maximum tolerated dose; rCFO, randomized CFO.

the significantly improved performance with larger sample sizes compared with smaller ones, indicating that the trial would converge to the MTD as the sample size increases. In scenario 5, where the true MTD is observed at the highest dose, no doses above the MTD are selected or allocated, resulting in 0 values for overdose selection and allocation. In scenario 2, where the lowest dose exceeds the target DLT by 0.1, early stopping occurs, and MTD selection is defined as the percentage of early stops, with no values for MTD allocation, overdose selection, and overdose allocation. These findings align with additional simulations on the basis of eight scenarios from the study by Yuan et al,[24] as shown in the Data Supplement. Even small improvements in the accuracy and safety can significantly affect cancer trials and new drug development, thus leading to life savings worldwide.

MTD Selection

When completing dose assignment in a real trial, an isotonic regression is conducted on the observed DLT rates to derive the final estimates using the pool-adjacent-violators algorithm.[22] The MTD is selected as the dose level whose isotonic estimate of the DLT rate is closest to the target rate ϕ. The function CFO.selectmtd() is used to select the MTD in single-drug trials, whereas CFO2d.selectmtd() is used for drug-combination trials (see the Data Supplement for details).

SIMULATION

Simulations help evaluate trial design performance by running a large number of simulated trials (multiple simulations). This section provides details on simulations using the CFO package.

Execution of One Simulation

In the context of simulated trials, the CFO package executes a single simulation of a CFO-type design using CFO.simu() for the CFO, aCFO, and rCFO designs; CFO2d.simu() for the 2dCFO design; and lateonset.simu() for designs with late-onset toxicity. The argument design is used to select different designs. Further implementation details are provided in the Data Supplement. Taking the f-aCFO design as an illustration, the following code executes the design, displaying the output in a textual summary and plotting the trajectory of dose level movements (Fig 4):

```
R > faCFOtrial <− lateonset.simu
(design = 'f − aCFO', target = 0.2, p.true
 = p.true, init.level = 1, ncohort = 12, cohortsize = 3,
assess.window = 3, tte.para = 0.5, accrual.rate = 2,
accrual.dist = 'unif', cutoff.eli = 0.95,
early.stop = 0.95, seed = 1)
R > summary(faCFOtrial)
R > plot(faCFOtrial)
```

Operating Characteristic Evaluation With Multiple Simulations

Extensive simulations are crucial for evaluating the operational characteristics of trial designs. The functions

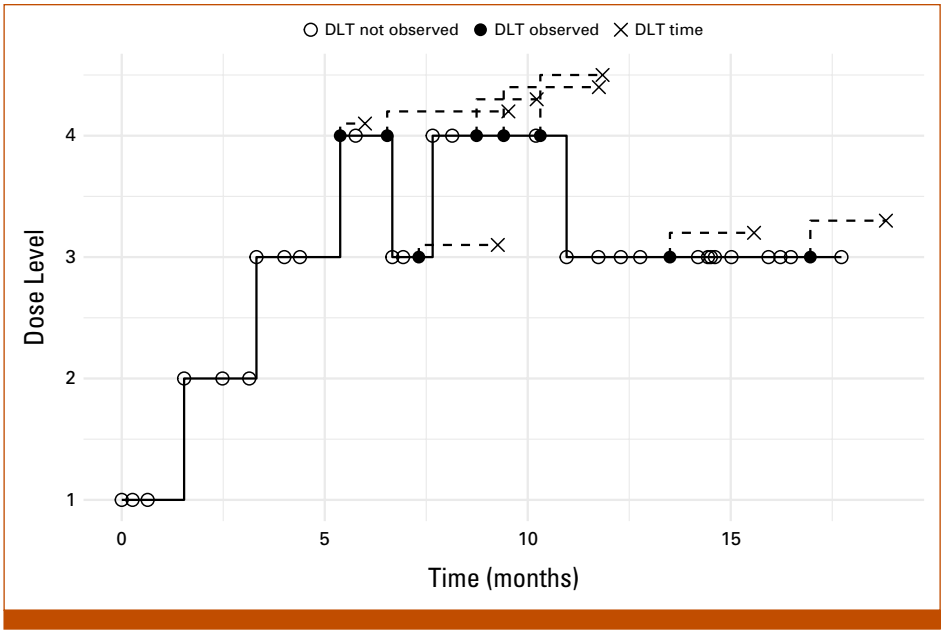

FIG 4. Illustration of a trial using the f-aCFO design. Patients are treated in a cohort of size 3, where solid circle ● and empty circle ○ indicate the presence or absence of observed toxicity in patients, respectively, and the x-axis value of the cross × signifies the time at which the DLT eventually occurred. CFO, calibration-free odds; DLT, dose-limiting toxicity; f-aCFO, fractional accumulative CFO.

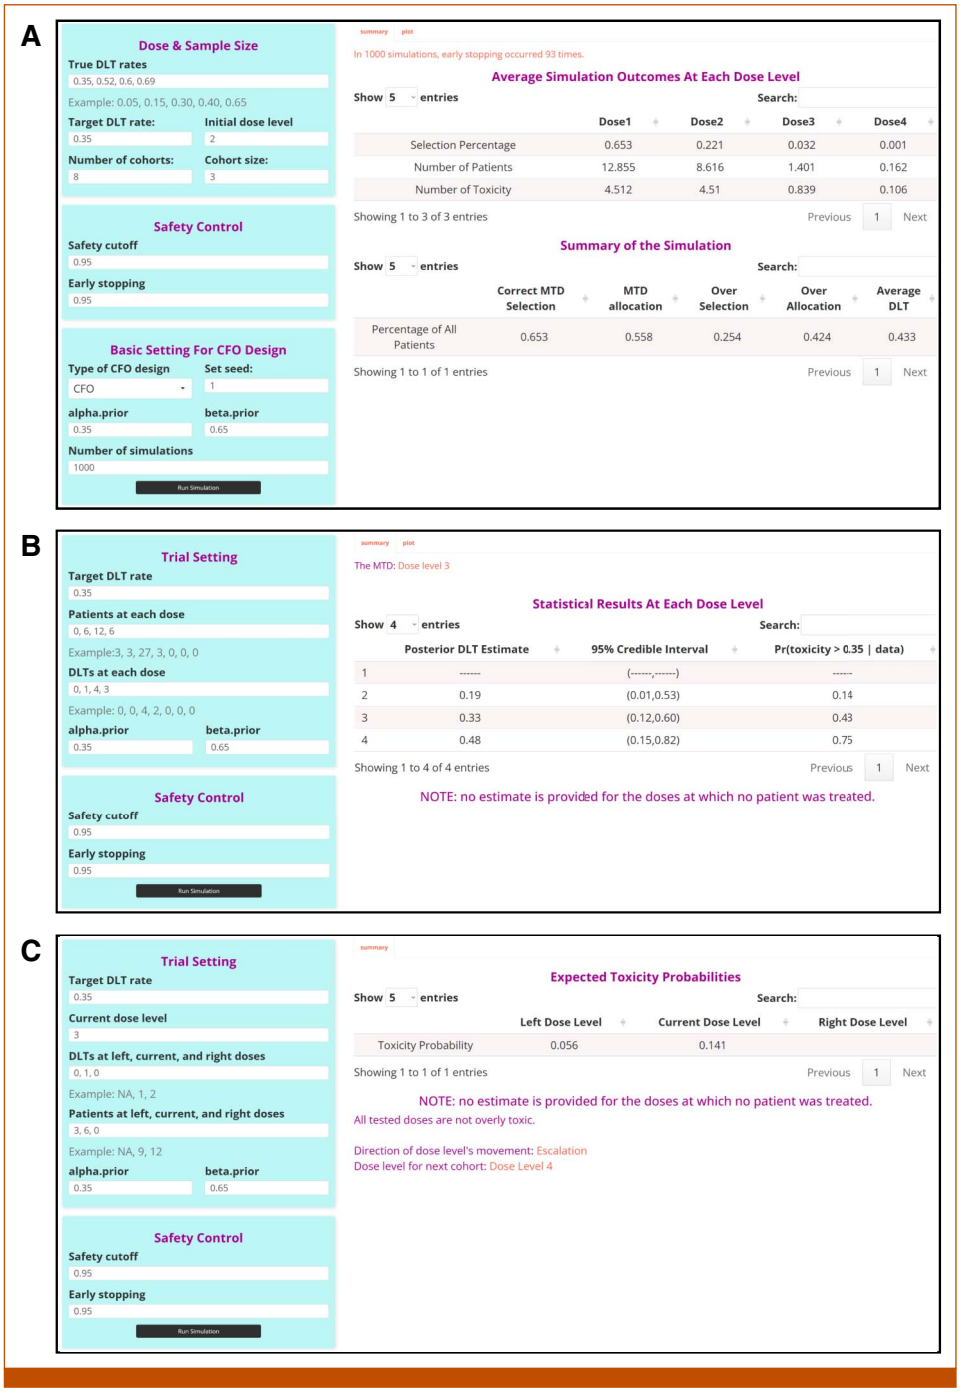

FIG 5. Use the CFO suite to conduct (A) Multiple Simulation, (B) Dose Level for Next Cohort, and (C) Select MTD for the trial example using the CFO design. CFO, calibration-free odds; MTD, maximum tolerated dose.

CFO.oc() and CFO2d.oc() facilitate multiple simulations for single-drug and drug-combination trials, respectively. For designs that do not involve late-onset toxicity (such as CFO, aCFO, and rCFO), time-related arguments like assess.window, accrual.rate, tte.para, and accrual.dist are set to NA. For designs with late-onset toxicity, such as TITE-CFO or f-aCFO, specific values are assigned to these arguments. The following code executes the

f-aCFO design, displaying the output in a textual summary:

```
R > faCFOoc <- CFO.oc(nsimu = 5000,
design = 'f - aCFO', target = 0.2, p.true = p.true,
init.level = 1, ncohort = 12, cohortsize = 3, assess.
window = 3, tte.para = 0.5, accrual.rate = 2, accrual.
dist = 'unif', prior.para = prior.para, cutoff.eli = 0.95,
early.stop = 0.95, seeds = 1 : 5000)
R > summary(faCFOoc)
R > plot(faCFOoc)
```

Comprehensive results and additional details on multiple simulations using CFO.oc() for single-drug trials and CFO2d.oc() for drug-combination trials are provided in the Data Supplement.

A REAL TRIAL EXAMPLE

This section illustrates the CFO design with a real trial example, using CFO suite to showcase its practical implementation.

The trial example is from a phase I-II study at the National Cancer Institute (ClinicalTrials.gov identifier: NCT02942264), aiming to determine the MTD of TG02, a pyrimidine-based multikinase inhibitor, in combination with temozolomide in adult patients with recurrent anaplastic astrocytoma or glioblastoma. The dose-finding phase aimed to enroll up to 24 patients in cohorts of size 3, with a target DLT probability of 0.35 because of a lack of effective treatments for this population. Four dose levels of TG02 (150, 200, 250, and 300 mg) are evaluated, with 200 mg selected as the starting dose on the basis of TG02's toxicity profile in other cancers. TG02 is administered in a 28-day cycle, with participants taking TG02 three days before the start of cycle 1 and then on four days during each cycle.

For safety, early stopping and dose elimination rules are adopted. A dose level k is deemed overly toxic if $\Pr(p_k > \phi \mid x_k, m_k \geq 3) > 0.95$. If the lowest dose level is overly toxic, indicated by $\Pr(p_1 > \phi \mid x_1, m_1 \geq 3) > 0.95$, the trial will be terminated according to the early stopping rule. Any dose level k identified as overly toxic, along with all the higher dose levels, is eliminated from further dose allocations to prioritize patient safety.

In our web application, the Multiple Simulation, Dose Level for Next Cohort, and Select MTD tabs are used sequentially for a comprehensive CFO design (Fig 5A-5C). To assess the precision and safety of the design before the actual trial, the Multiple Simulation tab is used to evaluate the probability of correctly identifying the true MTD and the distribution of patients across different dose levels. An effective design should exhibit a higher probability of correctly selecting the true MTD across various scenarios and allocate a substantial number of patients to this dose. Considering four scenarios with distinct true DLT rates: (0.35, 0.52, 0.60, 0.69), (0.16, 0.35, 0.50, 0.72), (0.02, 0.17, 0.35, 0.52), and (0.05, 0.12, 0.20, 0.35), where the MTDs are positioned at dose levels 1, 2, 3, and 4, respectively, the input settings and outcomes for scenario 1 are depicted in Figure 5A.

On commencing the trial, the first cohort is assigned to dose level 2. Subsequently, as each new cohort enrolls, the Dose Level for Next Cohort tab is used to input the outcomes of previously enrolled cohorts to determine the appropriate dose level for the upcoming cohort. For instance, if the current dose level is 3, and the cumulative numbers of DLTs and patients observed at left (dose level 2), current (dose level 3), and right (dose level 4) dose levels are (0, 1, 0) and (3, 6, 0), respectively, the next cohort will remain at dose level 3, as illustrated in Figure 5B. After assigning dose levels for all enrolled cohorts, the Select MTD tab is used to determine the MTD. Assume that at the end of the trial, the numbers of patients and DLTs at dose levels 2, 3, and 4 are (6, 12, 6) and (1, 4, 3), respectively, and no patients are treated at dose level 1, as shown in Figure 5C. Clicking on *Run Simulation* produces the estimated MTD (dose level 3), along with the estimated DLTs and their corresponding 95% credible intervals.

DISCUSSION

The CFO package is a user-friendly tool for implementing various CFO-type designs in phase I trials, covering key aspects such as dose determination, MTD selection, and simulation of operating characteristics. Its flexibility allows customization on the basis of dose information, late-onset toxicity, and single or combination therapies, making CFO applicable to diverse clinical settings. In addition, CFO provides intuitive descriptive and graphical outputs, enabling researchers to effectively communicate results to clinicians. Complementing the package, we developed CFO Suite, an interactive web application that inherits the CFO package's versatility and user-friendliness, providing a streamlined platform for applying CFO-type designs.

Together, CFO and CFO suite offer a comprehensive and accessible solution for optimizing phase I trial designs, bridging statistical theory and practical usability.

AFFILIATION

[1]School of Computing and Data Science, The University of Hong Kong, Hong Kong, China

PREPRINT VERSION

We disclose that a previous version of this manuscript was published as a preprint on medRxiv on July 27, 2024, with the following DOI: https://doi.org/10.1101/2024.07.26.24311051. The current manuscript has

been revised and updated based on the comments provided by two referees during the peer-review process.

CORRESPONDING AUTHOR

Guosheng Yin, PhD; e-mail: gyin@hku.hk.

SUPPORT

Supported in part by funding from the Research Grants Council of Hong Kong (17308321).

AUTHOR CONTRIBUTIONS

Conception and design: All authors
Collection and assembly of data: All authors
Data analysis and interpretation: All authors
Manuscript writing: All authors
Final approval of manuscript: All authors
Accountable for all aspects of the work: All authors

AUTHORS' DISCLOSURES OF POTENTIAL CONFLICTS OF INTEREST

The following represents disclosure information provided by authors of this manuscript. All relationships are considered compensated unless otherwise noted. Relationships are self-held unless noted. I = Immediate Family Member, Inst = My Institution. Relationships may not relate to the subject matter of this manuscript. For more information about ASCO's conflict of interest policy, please refer to www.asco.org/rwc or ascopubs.org/cci/author-center.
Open Payments is a public database containing information reported by companies about payments made to US-licensed physicians (Open Payments).

No potential conflicts of interest were reported.

ACKNOWLEDGMENT

We thank the two referees and Associate Editor for their many constructive and insightful comments that have led to significant improvements in the article.

REFERENCES

1. Yin G: Clinical Trial Design: Bayesian and Frequentist Adaptive Methods. Hoboken, NJ, John Wiley & Sons, 2012
2. Storer BE: Design and analysis of phase I clinical trials. Biometrics 45:925-937, 1989
3. Le Tourneau C, Lee JJ, Siu LL: Dose escalation methods in phase I cancer clinical trials. J Natl Cancer Inst 101:708-720, 2009
4. O'Quigley J, Pepe M, Fisher L: Continual reassessment method: A practical design for phase 1 clinical trials in cancer. Biometrics 46:33-48, 1990
5. Babb J, Rogatko A, Zacks S: Cancer phase I clinical trials: Efficient dose escalation with overdose control. Stat Med 17:1103-1120, 1998
6. Ivanova A, Flournoy N, Chung Y: Cumulative cohort design for dose-finding. J Stat Plan Inference 137:2316-2327, 2007
7. Liu S, Yuan Y: Bayesian optimal interval designs for phase I clinical trials. J R Stat Soc Ser C Appl Stat:507-523, 2015
8. Johnson VE: Uniformly most powerful Bayesian tests. Ann Stat 41:1716, 2013
9. Jin H, Yin G: CFO: Calibration-free odds design for phase I/II clinical trials. Stat Methods Med Res 31:1051-1066, 2022
10. Jin H, Yin G: Time-to-event calibration-free odds design: A robust efficient design for phase I trials with late-onset outcomes. Pharm Stat 22:773-783, 2023
11. Fang J, Yin G: Fractional accumulative calibration-free odds (f-aCFO) design for delayed toxicity in phase I clinical trials. Stat Med 22:3210-3226, 2024
12. Wang W, Jin H, Zhang YD, et al: Two-dimensional calibration-free odds design for phase I drug-combination trials. Front Oncol 13:1294258, 2023
13. Yan F, Zhang L, Zhou Y, et al: BOIN: an R package for designing single-agent and drug-combination dose-finding trials using Bayesian optimal interval designs. J Stat Softw 94:1-32, 2020
14. Zhu J, Zhang J, Takeda K: TITEgBOIN: Time-to-Event Dose-Finding Design for Multiple Toxicity Grades. Comprehensive R Archive Network (CRAN), 2023. https://cran.r-project.org/package=TITEgBOIN
15. Sweeting M, Mander A, Sabin T: Bcrm: Bayesian continual reassessment method designs for phase I dose-finding trials. J Stat Softw 54:1-26, 2013
16. Cheung K: dfcrm: Dose-Finding by the Continual Reassessment Method. Comprehensive R Archive Network (CRAN), 2019. https://cran.r-project.org/package=dfcrm
17. Diniz MA: ewoc: Escalation with Overdose Control. Comprehensive R Archive Network (CRAN), 2020. https://cran.r-project.org/package=ewoc
18. Blanchard MS: TEQR: Target Equivalence Range Design. CRAN, Comprehensive R Archive Network, 2016. https://cran.r-project.org/package=TEQR
19. R Core Team: R: A Language and Environment for Statistical Computing. Vienna, Austria, Foundation for Statistical Computing, 2013
20. Fang J, Zhang N, Wang W, et al: CFO: CFO-Type Designs in Phase I/II Clinical Trials. Comprehensive R Archive Network (CRAN), 2024. https://cran.r-project.org/package=CFO
21. Fang J, Zhang N, Wang W, et al: Calibration-Free Odds (CFO) Design For Phase I Trial. https://clinicaltrialdesign.shinyapps.io/cfoapp/
22. Bril G, Dykstra R, Pillers C, et al: Algorithm AS 206: Isotonic regression in two independent variables. J R Stat Soc Ser C Appl Stat 33:352-357, 1984
23. Cheung YK, Chappell R: Sequential designs for phase I clinical trials with late-onset toxicities. Biometrics 56:1177-1182, 2000
24. Yuan Y, Lin R, Li D, et al: Time-to-event Bayesian optimal interval design to accelerate phase I trials. Clin Cancer Res 24:4921-4930, 2018