







**MAIN PAPER**

# Decision rules for identifying combination therapies in open-entry, randomized controlled platform trials

Elias Laurin Meyer<sup>1</sup>  | Peter Mesenbrink<sup>2</sup>  | Cornelia Dunger-Baldauf<sup>3</sup>  |  
Ekkehard Glimm<sup>3,4</sup>  | Yuhan Li<sup>2</sup>  | Franz König<sup>1</sup>  | EU-PEARL (EU Patient-  
cEntric clinicAl tRial pLatforms) Consortium

<sup>1</sup>Center for Medical Statistics, Informatics, and Intelligent Systems, Medical University of Vienna, Vienna, Austria

<sup>2</sup>Analytics Department, Novartis Pharmaceuticals Corporation, East Hanover, New Jersey, USA

<sup>3</sup>Analytics Department, Novartis Pharma AG, Basel, Switzerland

<sup>4</sup>Institute of Biometry and Medical Informatics, University of Magdeburg, Magdeburg, Germany

**Correspondence**

Franz König, Center for Medical Statistics, Informatics, and Intelligent Systems, Medical University of Vienna, Vienna, Austria.  
Email: [franz.koenig@meduniwien.ac.at](mailto:franz.koenig@meduniwien.ac.at)

**Funding information**

EU-PEARL (EU Patient-cEntric clinicAl tRial pLatforms) project has received funding from the Innovative Medicines Initiative (IMI) 2 Joint Undertaking (JU) under grant agreement No 853966. This Joint Undertaking receives support from the European Union's Horizon 2020 research and innovation programme and EFPIA and Children's Tumor Foundation, Global Alliance for TB Drug Development non-profit organization, Springworks Therapeutics Inc. This publication reflects the authors' views. Neither IMI nor the European Union, EFPIA, or any Associated Partners are responsible for any use that may be made of the information contained herein. The PhD research of Elias Laurin Meyer was funded until 11/2020 by Novartis through the University and not at an individual level.

**Abstract**

Platform trials have become increasingly popular for drug development programs, attracting interest from statisticians, clinicians and regulatory agencies. Many statistical questions related to designing platform trials—such as the impact of decision rules, sharing of information across cohorts, and allocation ratios on operating characteristics and error rates—remain unanswered. In many platform trials, the definition of error rates is not straightforward as classical error rate concepts are not applicable. For an open-entry, exploratory platform trial design comparing combination therapies to the respective monotherapies and standard-of-care, we define a set of error rates and operating characteristics and then use these to compare a set of design parameters under a range of simulation assumptions. When setting up the simulations, we aimed for realistic trial trajectories, such that for example, a priori we do not know the exact number of treatments that will be included over time in a specific simulation run as this follows a stochastic mechanism. Our results indicate that the method of data sharing, exact specification of decision rules and a priori assumptions regarding the treatment efficacy all strongly contribute to the operating characteristics of the platform trial. Furthermore, different operating characteristics might be of importance to different stakeholders. Together with the potential flexibility and complexity of a platform trial, which also impact the achieved operating characteristics via, for example, the degree of efficiency of data sharing this implies that utmost care needs to be given to evaluation of different assumptions and design parameters at the design stage.

**KEYWORDS**

adaptive trials, clinical trial simulation, combination therapy, platform trials

## 1 | INTRODUCTION

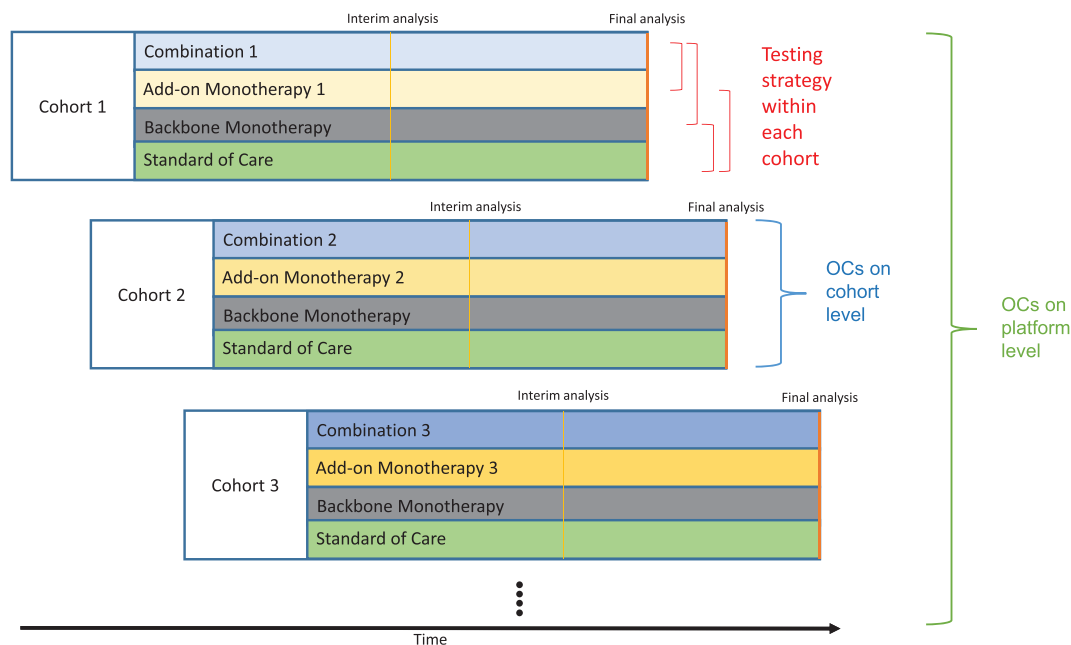
The goal to test as many investigational treatments as possible over the shortest duration, which is influenced by both recent advances in drug discovery and biotechnology and the ongoing global pandemic due to the SARS-CoV-2 virus,<sup>1-6</sup> has made master protocols and especially platform trials an increasingly feasible alternative solution to the time-consuming sequence of classical randomized controlled trials.<sup>7-12</sup> Platform trial designs allow for the evaluation of one or more investigational treatments in the study population(s) of interest within the same clinical trial, as compared to traditional randomized controlled trials, which usually evaluate only one investigational treatment in one study population. When cohorts share key inclusion/exclusion criteria, trial data can easily be shared across such sub-studies. In practice, setting up a platform trial potentially may require additional time due to operational and statistical challenges. However, simulations have shown that platform trials can be superior to classical trial designs with respect to various operating characteristics which include the overall study duration. In a setting where only few new agents are expected to be superior to standard of care, Saville and Berry<sup>13</sup> investigated the operating characteristics of adaptive Bayesian platform trials using binary endpoints compared with a sequence of “traditional” trials, that is, trials testing only one hypothesis, and found that platform trials perform dramatically better in terms of the number of patients and time required until the first superior investigational treatment has been identified. Using real data from the 2013 to 2016 Ebola virus disease epidemic in West Africa, Brueckner et al.<sup>14</sup> investigated the operating characteristics of various multi-arm multi-stage and two-arm single stage designs, as well as group-sequential two-arm designs, and found that designs with frequent interim analyses outperformed single-stage designs with respect to average trial duration and sample size when fixing the type 1 error and power. When having a pool of investigational agents available, which should all be tested against a common shared control, and using progression-free survival as the efficacy endpoint, Yuan et al.<sup>15</sup> found that average trial duration and average sample size are drastically reduced when using a multi-arm, Bayesian adaptive platform trial design using response-adaptive randomization compared with traditional two-arm trials evaluating one agent at a time. Hobbs et al.<sup>16</sup> reached a similar conclusion when comparing a platform trial with a binary endpoint and futility monitoring based on Bayesian predictive probabilities with a sequence of two-arm trials. Tang et al.<sup>17</sup> investigated a phase II setting in which several monotherapies are combined with several backbone therapies and tested in a single-arm manner. Assuming different treatment combination effects, they found that their proposed Bayesian platform design with adaptive shrinkage has a lower average sample size in the majority of scenarios investigated and always a higher percentage of correct combination selections when compared with a fully Bayesian hierarchical model and a sequence of Simon's two-stage designs. Ventz et al.<sup>18</sup> proposed a frequentist adaptive platform (so called “rolling-arms design”) design as an alternative to sequences of two-arm designs and Bayesian adaptive platform designs, which is much simpler than the Bayesian adaptive platform designs in that it uses equal allocation ratios and simpler and established decision rules based on group sequential analysis. The authors found that performance under different treatment effect assumptions and a set of general assumptions was comparable to, if not slightly better than Bayesian adaptive platform designs and much better than a sequence of traditional two-arm designs in terms of average sample size and study duration. For a comprehensive review on the evolution of master protocol clinical trials and the differences between basket, umbrella and platform trials, see Meyer et al.<sup>11</sup> In this article, we explore the impact of both decision rules and assumptions on the nature of the treatment effects and availability of treatments on certain operating characteristics of an open-entry, cohort platform trial with multiple study arms. Such a platform design makes sense in any context in which multiple pair-wise comparisons are necessary to advance a compound, whereby some comparisons are based on treatments unique to different cohorts and some comparisons are based on treatments common to all cohorts (and any mixture thereof). An example for such a context would be a drug development program in Nonalcoholic steatohepatitis (NASH) trying to advance two-compound combination therapies with a common backbone therapy.

The article is organized as follows: In Section 2, we describe the trial design under investigation, the different testing strategies as well as the investigated operating characteristics. In Section 3, we discuss the simulation setup and present and discuss the results of the different simulation scenarios. We conclude with a general discussion in Section 4.

## 2 | METHODS

### 2.1 | Platform design

We investigated an open-entry, exploratory cohort platform study design with a binary endpoint evaluating the efficacy of a two-compound combination therapy compared to the respective monotherapies and the standard-of-care (SoC).



**FIGURE 1** Schematic overview of the proposed platform trial design. New cohorts consisting of a combination therapy arm, a monotherapy arm using the same compound in every cohort (referred to as “backbone monotherapy”), an add-on monotherapy arm which is different in every cohort and a SoC arm are entering the platform over time. While the add-on monotherapy and therefore the combination therapy is different in every cohort (as indicated by the differently shaded colors), the backbone monotherapy and SoC are the same in every cohort (as indicated by the same colors). Each cohort has an interim analysis after about half of the initially planned sample size, after which the cohort can be stopped for early efficacy or futility. The red brackets indicate the testing strategy within each cohort, that is, comparison of combination therapy against both monotherapies and both monotherapies against SoC. We differentiate between per-cohort and per-platform operating characteristics (OCs)

After an initial inclusion of one cohort, we allow new cohorts to enter the platform trial over time until a maximum number of cohorts is reached. We assume different cohorts share common key inclusion/exclusion criteria and patients are always drawn from the same population, such that all eligible patients can be randomized to any cohort. Each cohort consists of four arms: combination therapy, monotherapy A, monotherapy B and SoC. Monotherapy A is the same for all cohorts (further referred to as “backbone monotherapy”), while monotherapy B (further referred to as “add-on monotherapy X”) is different in every cohort X. The combination of monotherapies A and B is called *combination therapy*. See Figure 1 for a schematic overview of the proposed trial design. To show that the backbone and SoC treatments do not change during the platform trial, we use the same color coding of gray and green for all cohorts, respectively. The x-axis shows the calendar time. At any point in time, new cohorts could enter the platform trial. In contrast to a simulation of a classical RCT, here different response rates for the same treatment arms in different cohorts can be assumed. The assignment of response rates to treatment arms is described in more detail in Appendix A.3. We conduct one interim analysis for every cohort (indicated by the vertical yellow line in Figure 1), on the basis of which the cohort might be stopped early for either futility or efficacy. The platform trial ends if there is no active, recruiting cohort left. The proposed trial design is suggested for phase II trials. Subsequent phase III studies would follow if a combination therapy is graduated successfully. Following a rigorous interpretation of the current FDA and EMA regulatory guidelines,<sup>19,20</sup> superiority of the combination therapy over both monotherapies and superiority of both monotherapies over SoC needs to be demonstrated. Therefore, we test the combination therapy against both of the monotherapies and both of the monotherapies against SoC, resulting in four comparisons (indicated by the red “testing strategy” lines in Figure 1).

## 2.2 | Bayesian decision rules

As we discussed in Section 2.1, for testing the efficacy of the combination therapy, four pair-wise comparisons are performed. In this article, for every one of the comparisons, we consider Bayesian decision rules based on the posterior

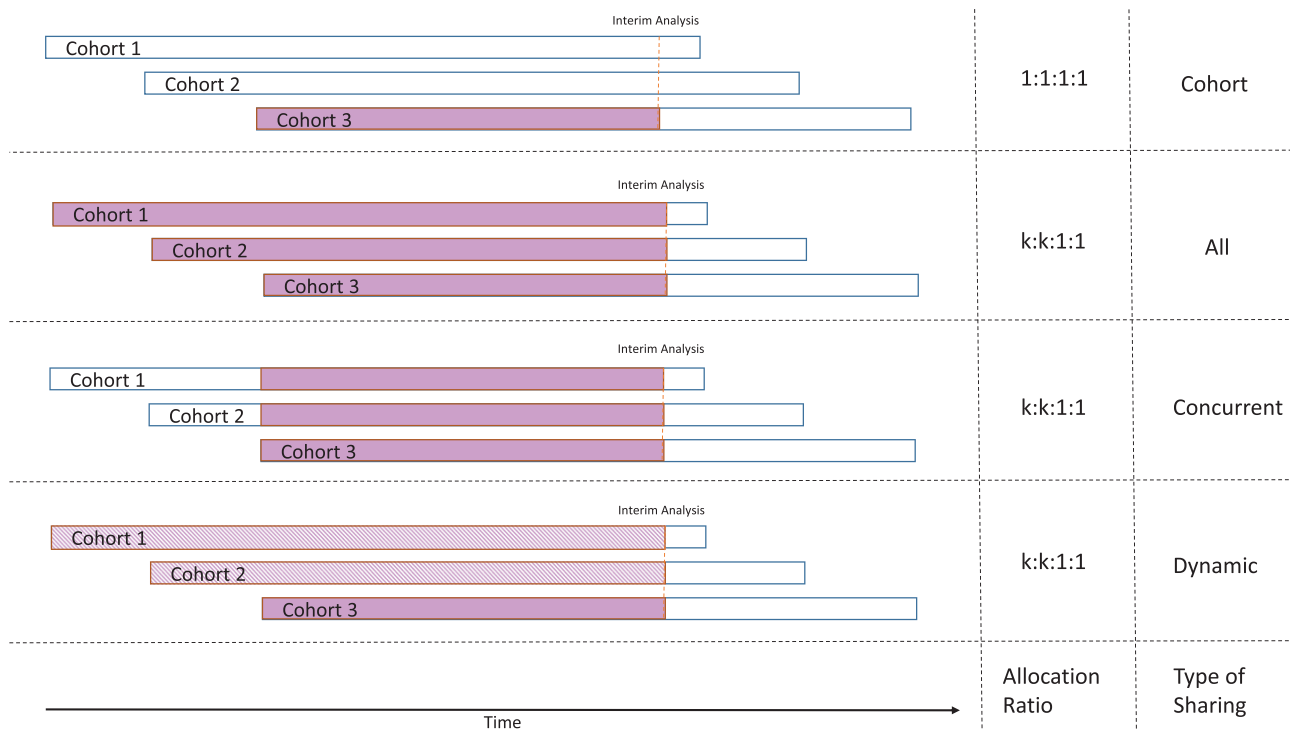
distributions of the response rates of the respective study arms.<sup>21</sup> In principle, any other Bayesian (e.g., predictive probabilities, hierarchical models, etc.) or frequentist rules (based on  $p$ -values, point estimates or confidence intervals) could be used. While these decision rules are based on fundamentally different paradigms, they might translate into the exact same stopping rules, for example, with respect to the observed response rate.<sup>22,23</sup> Decision rules based on posterior distributions used vague, independent Beta(1/2, 1/2) priors for all simulation results presented in this article; however, please note that independence of the vague priors is a strong assumption and its violation can lead to selection bias as pointed out by many authors.<sup>24–26</sup> We differentiate two types of decisions for cohorts: “GO” (graduate combination therapy, that is, declare combination therapy successful) and “STOP” (stop evaluation of combination therapy and do not graduate, that is, declare the combination therapy unsuccessful). Generally, we consider decision rules of the following form:

$$\begin{aligned}
 \text{GO, if } & (P(\pi_C > \pi_A + \delta_{CA}^{E,T} | \text{Data}) > \gamma_{CA}^{E,T}) \wedge \\
 & (P(\pi_C > \pi_B + \delta_{CB}^{E,T} | \text{Data}) > \gamma_{CB}^{E,T}) \wedge \\
 & (P(\pi_A > \pi_S + \delta_{AS}^{E,T} | \text{Data}) > \gamma_{AS}^{E,T}) \wedge \\
 & (P(\pi_B > \pi_S + \delta_{BS}^{E,T} | \text{Data}) > \gamma_{BS}^{E,T}) \\
 \\
 \text{STOP, if } & (P(\pi_C > \pi_A + \delta_{CA}^{F,T} | \text{Data}) < \gamma_{CA}^{F,T}) \vee \\
 & (P(\pi_C > \pi_B + \delta_{CB}^{F,T} | \text{Data}) < \gamma_{CB}^{F,T}) \vee \\
 & (P(\pi_A > \pi_S + \delta_{AS}^{F,T} | \text{Data}) < \gamma_{AS}^{F,T}) \vee \\
 & (P(\pi_B > \pi_S + \delta_{BS}^{F,T} | \text{Data}) < \gamma_{BS}^{F,T})
 \end{aligned} \tag{1}$$

whereby  $\pi_x$  denotes the response rate in treatment arm  $x$  ( $x \in \{C, A, B, S\}$ ) and  $T \in 1, 2$  denotes the analysis time point. At interim ( $T = 1$ ), if neither a decision for early efficacy or futility is made, the cohort continues unchanged. At final ( $T = 2$ ), if the efficacy boundaries are not met, the cohort automatically stops for futility. The initial letters E or F in the superscript of the thresholds  $\delta$  and  $\gamma$  indicate if this boundary is used to stop for efficacy (E) or futility (F). The subscripts refer to the treatment arms (C (combination), A (monotherapy A), B (monotherapy B) and S (SoC)) and allow for different thresholds for the individual comparisons. Choosing, for example,  $\gamma_{CA}^{E,1} = \gamma_{CB}^{E,1} = \gamma_{AS}^{E,1} = \gamma_{BS}^{E,1} = 1$  corresponds to not allowing early stopping for efficacy at interim. If at interim both stopping for early efficacy and futility is allowed, parameters need to be chosen carefully such that GO and STOP and decisions are not simultaneously possible.

### 2.3 | Data sharing and allocation ratios

The advantage of platform trials is that they can use the data from all cohorts, which potentially increases the power in every individual cohort. We can specify whether we want to share information on the backbone monotherapy and SoC arms across the study cohorts. Several different methods have been proposed to facilitate adequate borrowing of non-concurrent (these can be internal or external to the trial) controls.<sup>27–32</sup> We consider four options, all applying to both SoC and backbone monotherapy: (1) no sharing, using only data from the current cohort (see first row of Figure 2), (2) full sharing of all available data, that is, using all data 1-to-1 (see second row of Figure 2), (3) only sharing of concurrent data, that is, using concurrent data 1-to-1 (see third row of Figure 2), and (4) using a dynamic borrowing approach further described in Appendix A.1, in which the degree of shared data increases with the homogeneity of the observed data in the current cohort and the pooled observed data of all other cohorts, that is, discounting the pooled observed data of other cohorts less, if the observed response rate is similar (see fourth row of Figure 2). Whenever a new cohort enters the platform trial, a key design element is the allocation ratio to the newly added arms (combination therapy, add-on monotherapy, backbone monotherapy and SoC) and whether the allocation ratio of the already ongoing cohorts should be changed as well, for example, randomizing less patients to backbone monotherapy and SoC in case this data is shared across cohorts. The platform trial advances dynamically and as a result the structure can follow many different trajectories (it is unknown in the beginning of the trial how many cohorts will enter the platform, how many of them will run concurrently, whether the generated data will stem from the same underlying distributions and should therefore be shared, etc.). As the best possible compromise under uncertainty, we aimed to achieve a balanced randomization for every comparison in case of either no data sharing or sharing only concurrent data. Depending on the type of



**FIGURE 2** Schematic overview of the different levels of sharing. No sharing happens if only “cohort” data are used. If sharing “all” data, whenever in any cohort an interim or final analysis is performed, all SoC and backbone monotherapy data available from all cohorts are used. If sharing only “concurrent” data, whenever in any cohort an interim or final analysis is performed, all SoC and backbone monotherapy data that was collected during the active enrollment time of the cohort under investigation are used. If sharing “dynamically,” whenever in any cohort an interim or final analysis is performed, the degree of data sharing of SoC and backbone monotherapy data from other cohorts increases with the homogeneity of the observed response rate of the respective arms. A solid fill represents using data 1-to-1, while a dashed fill represents using discounted data (for more information see Appendix A.1). If at any given time there are  $k$  active cohorts, the allocation ratio is 1:1:1:1 in case of no data sharing and  $k:k:1:1$  otherwise (combination: add-on monotherapy: backbone monotherapy: SoC). This allocation ratio is updated for all active cohorts every time the number of active cohorts  $k$  changes due to dropping or adding a new cohort

data sharing and the number of active arms, this means either a balanced randomization ratio within each cohort in case of no data sharing (i.e., 1:1:1:1, combination: add-on mono (monotherapy B): backbone mono (monotherapy A): SoC), or a randomization ratio that allocates more patients to the combination and add-on monotherapy arm for every additional active cohort in case of using only concurrent data. As an example, if at any point in time  $k$  cohorts are active at the same time and we share only concurrent data, the randomization ratio is  $k:k:1:1$  in all cohorts, which ensures an equal number of patients per arm for every comparison. This allocation ratio is updated for all active cohorts every time the number of active cohorts  $k$  changes due to dropping or adding a new cohort. As an example, assume 30:30:30:30 patients have been enrolled in cohort 1 before a second cohort is added. Then, until, for example, an interim analysis is performed in cohort 1, both cohorts will have an allocation ratio of 2:2:1:1. If the interim analysis is scheduled after 180 patients per cohort, this would mean another 20:20:10:10 patients need to be enrolled in cohorts 1 and 2, since we can use the 10 concurrently enrolled backbone monotherapy and 10 concurrently enrolled SoC patients in cohort 2 for the interim analysis in cohort 1, leading to a balanced 50:50:50:50 patients for the comparisons. In case of either full sharing or dynamic borrowing, we use the same approach as when using concurrent data only.

## 2.4 | Definition of cohort success and operating characteristics

As discussed in section 2.1, for testing the efficacy of the combination therapy, four pair-wise comparisons are conducted. If we were running a single, independent trial investigating a combination therapy, we would consider the trial a success if all of the necessary pair-wise comparison were successful. Consequently, we would consider it a failure if at

least one of the necessary pair-wise comparison was unsuccessful. For clearer understanding, we call these two options, respectively, a positive or negative outcome. Depending on the formulated hypotheses, these might be either true or false positives or negatives. To allow evaluation of the decision rules in terms of frequentist type 1 error and power, for each of the four pair-wise comparisons we formulate a set of hypotheses of the following sort (exemplary for combination vs. monotherapy A, but analogously for all other comparisons):  $H_0: \pi_C \leq \pi_A + \zeta_{CA}$  versus  $H_1: \pi_C > \pi_A + \zeta_{CA}$ . Please note that as we are drawing the response rates randomly for every new cohort entering the platform trial, it is unknown a priori whether this null hypothesis holds. Furthermore, please note while for our simulations this is not the case,  $\zeta_{CA}$  can be different from  $\delta_{CA}$  chosen in Section 2.2 (in our simulations, for all pair-wise comparisons  $\zeta = \delta = 0$ ). We will elaborate more on what happens in this case in Appendix A.2. If all such pair-wise alternative hypotheses within a cohort hold, we call the cohort truly efficacious (and any decision made is either a true positive or a false negative), otherwise we call it truly not efficacious (and any decision made is either a true negative or a false positive). The latter implies that this includes scenarios where for some of the pair-wise hypotheses of interest the alternative and for some the null holds. Depending on the decision rules used, it could be debated if other definitions are more meaningful (e.g., if the comparisons of the combination therapy to the monotherapies are both true alternatives the cohort may be considered as efficacious). While for a single, independent trial this would yield a single outcome (true positive, false positive, true negative, false negative), for a platform trial with multiple cohorts this yields a vector of such outcomes, one for each investigated cohort. In order to evaluate different trial designs, operating characteristics need to be chosen that take into account the special features of the trial design, but are at the same time interpretable in the classical context of hypothesis testing. In more simple trial designs evaluating one treatment against a control, power and type 1 error rates are used to judge the design under consideration. For platform trials, similarly to multi-arm multi-stage trials,<sup>33</sup> the choice of operating characteristics is not obvious.<sup>34–38</sup> Furthermore, for the particular trial design under consideration, many operating characteristics based on the pair-wise comparisons of the different monotherapies, SoC and combination therapy could be considered. We decided to investigate cohort-level and platform-level operating characteristics (see Figure 1). As described previously, when simulating platform trials we might allow the same treatment arms to have different response rates in different cohorts. Even if we allow for a mix of true null and alternative hypotheses, for a single simulated platform trial, it could happen by chance that only efficacious or only not efficacious cohorts are added. Consequently, for the platform-level operating characteristics, a definition challenge arises when by chance either no true positive and false negative (in case all cohorts are truly not efficacious) or true negative and false positive (in case all cohorts are truly efficacious) decisions are possible. To make sure that the operating characteristics reflect this situation (which depends on the prior on the treatment effects), we differentiate between counting all simulation iterations (which implicitly takes into account the prior on the treatment effect, “BA” [“Bayesian Average”] operating characteristics, FWER BA and Disj Power BA) or only those simulation iterations where a false decision could have been made towards the type 1 error rate and power (FWER and Disj Power). An overview of the operating characteristics used in this article and their definitions can be found in Table 1.

### 3 | SIMULATIONS

#### 3.1 | Simulation setup

We investigate the impact of a range of different design parameters and assumptions on the operating characteristics. In total, we investigated 14 different settings with respect to the treatment efficacy assumptions for the combination arm, the monotherapy arms and the SoC arm. In the main text, with one exception, only results of one treatment efficacy setting (setting 1) are shown as in this scenario both truly efficacious and not efficacious cohorts can enter the platform trial (see section 2.4). In the investigated setting (setting 1), for every new cohort entering the platform trial, the backbone monotherapy (response rate 20%) is superior to SoC (response rate 10%). The add-on monotherapy efficacy is random with 50% probability to be as efficacious as the backbone monotherapy (response rate 20%) and 50% probability to be not efficacious (response rate 10%). Adding to the monotherapies, the combination therapy interaction effect is additive, meaning the combination therapy is superior to both monotherapies if the add-on monotherapy is efficacious (response rate 40%) and not superior to the backbone monotherapy otherwise (response rate 20%). In terms of sample sizes, we vary the final cohort sample size from 100 to 500 in steps of 100 and fix the interim sample size at half of the final sample size. The maximum number of cohorts entering the platform trial over time is varied between 3 and 7 in steps of 2 and the probability to include new cohorts in the platform after every patient is set to either 1% or 3%. For the

TABLE 1 Operating characteristics used in this paper and their definitions

Name	Definition
PCP	“Per-Cohort-Power,” the ratio of the sum of true positives among the sum of truly efficacious cohorts (i.e., the sum of true positives and false negatives) across all platform trial simulations, that is, the probability of a true positive decision for any new cohort entering the trial. This is a measure of how wasteful the trial is with superior therapies.
PCT1ER	“Per-Cohort-Type-1-Error,” the ratio of the sum of false positives among the sum of all truly not efficacious cohorts across all platform trial simulations, that is, the probability of a false positive decision for any new cohort entering the trial. This is a measure of how sensitive the trial is in detecting futile therapies.
FWER	“Family-wise type 1 Error Rate”, the proportion of platform trials, in which at least one false positive decision has been made (i.e., probability of at least one false positive decision across all cohorts), where only such trials are considered, which contain at least one cohort that is in truth futile. Formal definition: $\frac{1}{ I_0 } \sum_{i \in I_0} \mathbb{1}\{FP_i > 0\}$ , where $I_0^* = \{i \in \{1, \dots, iter\} : n_i^{H0} > 0\}$ , where <i>iter</i> is the number of platform trial simulation iterations, $FP_i$ denotes the number of false-positive decisions in simulated platform trial <i>i</i> and $n_i^{H0}$ is the number of not efficacious cohorts in platform trial <i>i</i> .
FWER BA	“Family-wise type 1 Error Rate Bayesian Average”, the proportion of platform trials, in which at least one false positive decision has been made (i.e., probability of at least one false positive decision across all cohorts), regardless of whether or not any cohorts which are in truth futile exist in these trials. Formal definition: $\frac{1}{iter} \sum_{i=1}^{iter} \mathbb{1}\{FP_i > 0\}$ , where <i>iter</i> is the number of platform trial simulation iterations and $FP_i$ denotes the number of false-positive decisions in simulated platform trial <i>i</i> . This will differ from FWER in scenarios where—due to a prior on the treatment effect—in some simulation runs, there are by chance only efficacious cohorts in the platform trial (see Appendix A.3 for more details on the different treatment efficacy scenarios).
Disj Power	“Disjunctive Power”, the proportion of platform trials, in which at least one correct positive decision has been made (i.e., probability of at least one true positive decision across all cohorts), where only such trials are considered, which contain at least one cohort that is in truth superior. Formal definition: $\frac{1}{ I_1 } \sum_{i \in I_1} \mathbb{1}\{TP_i > 0\}$ , where $I_1^* = \{i \in \{1, \dots, iter\} : n_i^{H1} > 0\}$ , where <i>iter</i> is the number of platform trial simulation iterations, $TP_i$ denotes the number of true-positive decisions in simulated platform trial <i>i</i> and $n_i^{H1}$ is the number of efficacious cohorts in platform trial <i>i</i> .
Disj Power BA	“Disjunctive Power Bayesian Average”, the proportion of platform trials, in which at least one correct positive decision has been made (i.e., probability of at least one true positive decision across all cohorts), regardless of whether or not any cohorts which are in truth superior exist in these trials. Formal definition: $\frac{1}{iter} \sum_{i=1}^{iter} \mathbb{1}\{TP_i > 0\}$ , where <i>iter</i> is the number of platform trial simulation iterations and $TP_i$ denotes the number of true-positive decisions in simulated platform trial <i>i</i> . This will differ from Disj Power in scenarios where—due to a prior on the treatment effect—in some simulation runs, there are by chance no efficacious cohorts in the platform trial (see Appendix A.3 for more details on the different treatment efficacy scenarios).

Bayesian decision rules (see section 2.2), by default we set all  $\delta = 0$ , all  $\gamma^{E,T} = 0.9$  and all  $\gamma^{F,T} = 0.5$  in Equation (1). An overview of the simulation setup can be found in Table 2. It should be obvious that the chosen decision rules will yield dramatically different power and type 1 errors across different simulation parameters and treatment efficacy settings. It was our primary goal to investigate the relative impact of the simulation parameters and treatment efficacy settings on the operating characteristics. Of course, for any given combination of simulation parameters and treatment efficacy assumptions, the decision rules could be adapted to, for example, achieve a per-cohort power of 80%. We will investigate the impact of the decision rules in more detail in section 3.2.1.

For every configuration of design parameters and assumptions, 10,000 platform trials were simulated. Unless otherwise specified, simulation results presented in this article use treatment effect scenario 1, the above mentioned Bayesian decision rules, a final sample size of 500 per cohort, a maximum number of cohorts per platform trial of up to 7 and a probability of including a new cohort after every patient of 3%. Results of all further treatment efficacy assumptions, as well as a table summarizing all different treatment efficacy settings are presented in the supplements (Table 3). For a detailed overview of the general simulation assumptions, as well as all possible simulation parameters for the R software package and Shiny App, see the R package **CohortPlat** vignette on GitHub. The R software package can be downloaded from GitHub or CRAN. Further results not discussed in the main text can be found in the Supplemental material. For a complete overview of all simulation results, we developed a Shiny App facilitating self-exploration of all of our simulation results. The purpose of this R Shiny app is to quickly inspect and visualize all simulation results that were computed for this paper. We uploaded the R Shiny app, alongside all of our simulation results used in this paper to our server.<sup>39</sup>

**TABLE 2** Simulation Setup Overview. For different simulation parameters, we differentiate between parameters that are considered a design choice and parameters that are considered an assumption regarding the future course of the platform or treatment effects. For the investigated values, we state in bold which value was considered the default value, that is, unless stated otherwise in a particular figure, the parameters was set to this value. For some parameters there is no default (e.g., when shown as a simulation dimension in every figure or if chosen as a fixed design parameter in section 2)

Name	Type	Investigated values	Description
Maximum number of cohorts	Assumption	3, 5, 7	Assumed maximum number of cohorts per platform (can be less in individual simulations)
Cohort inclusion rate	Assumption	0.01, <b>0.03</b>	Probability to include a new cohort in the ongoing platform trial after every simulated patient (unless maximum number of cohorts is reached). The default value leads to reaching the assumed maximum number of cohorts in nearly every simulation for nearly every sample size.
Treatment efficacy setting	Assumption	<b>1</b> , 2–14	Assumed treatment effects of the different study arms. For more information, see Table 3.
Final cohort sample size	Design choice	100, 200, 300, 400, <b>500</b>	Number of patients after which final analysis in a cohort is conducted (interim analysis always after half the final sample size)
Data sharing	Design choice	All, concurrent, dynamic, cohort	Different methods of data sharing used at analyses, ranging from full pooling to not sharing at all. For more information, see section 2.3.
Allocation ratios	Design choice	Balanced or unbalanced (depends on data sharing)	Allocation ratio to treatment arms within a cohort. Depending on the method of data sharing, the allocation ratio is set to either balanced or unbalanced (i.e., randomizing more patients to combination and add-on monotherapy). For more information, see section 2.3.
Bayesian decision rule	Design choice	$\delta \in [0, 0.2]$ , $\gamma \in [0.65, 0.95]$ ( <b>0.9</b> )	Thresholds used in the Bayesian decision making at interim and final. Values other than the default values are used only in Figure 6. For more information, see section 2.2.
Interim analysis	Design choice	<b>Early stopping for futility</b> , no early stopping for futility	Binding rules on whether or not a cohort can be stopped at interim for early futility (note: cohorts can always stop for early efficacy). For more information, see section 2.2. Further extensions investigated in the supplements include using a surrogate short-term endpoint for interim decision making.

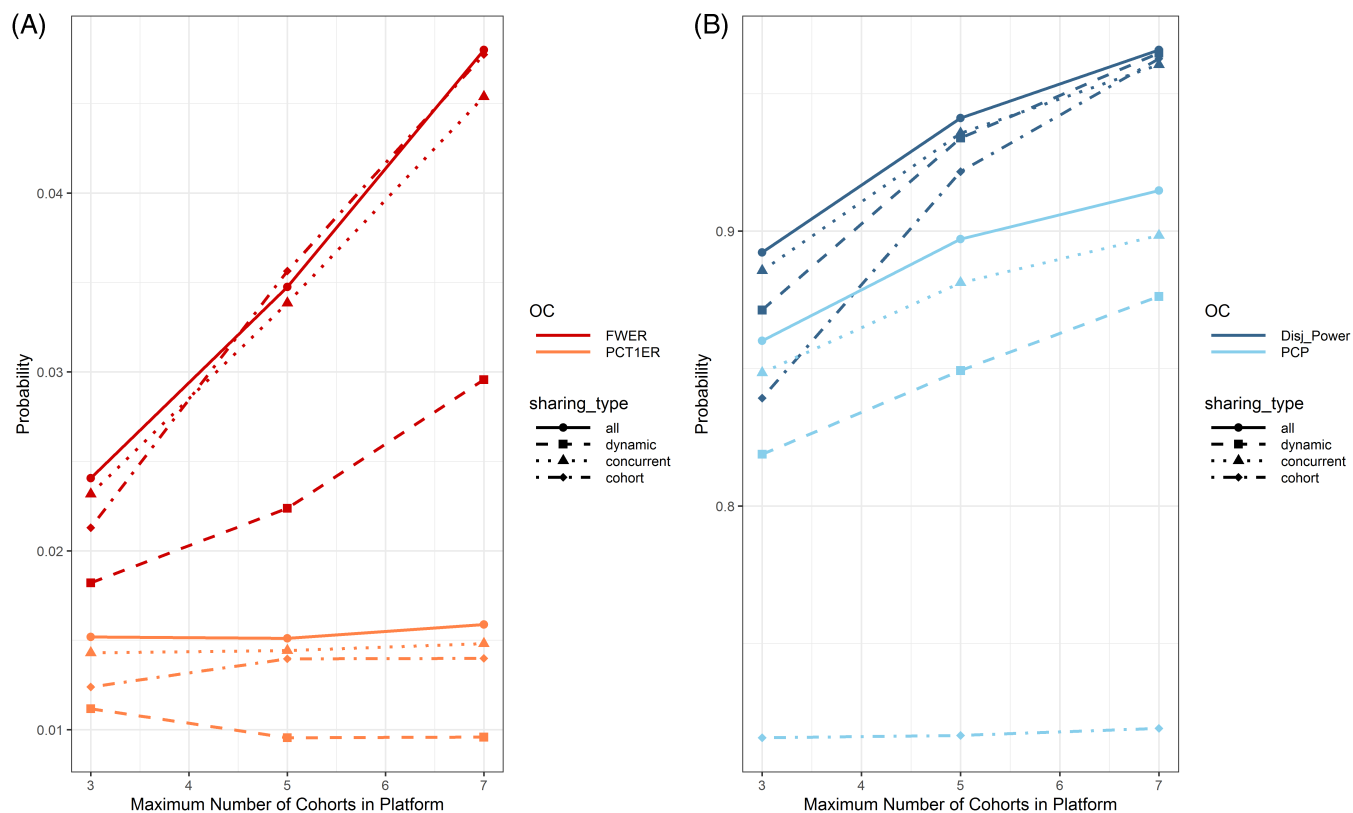
### 3.2 | Simulation results

In Figure 3, we investigate the impact of the data sharing and maximum number of cohorts per platform on PCT1ER and FWER (Figure 3A) and PCP and disjunctive power (Figure 3B). We observe that the PCP does not increase with the platform size (maximum number of cohorts) in case of no data sharing, while it does increase with the magnitude of data sharing. For the disjunctive power we observe a similar relationship, that is, it increases regardless of the data sharing with increasing number of investigated cohorts. Furthermore, we see that the PCT1ER stays more or less constant with respect to the platform size, however is lowest for the dynamic borrowing approach. We believe that this is due to the following: A type 1 error in a cohort happens only if either the observed combination therapy and add-on monotherapy efficacy are on a random high or if the observed backbone and SoC response rates are on a random low. In the former case, the degree of data sharing will have no impact on the decisions made. In the latter case, if dynamic borrowing is used, all other cohorts will discount this data, leading to fewer type 1 errors in the other cohorts compared to if complete pooling was used. On the other hand, for the cohort that had observed backbone and SoC response rates on a random low, there is a chance to not make a type 1 error when any sort of data sharing is used. As a result of the PCT1ER behavior, the FWER increases with increasing platform size and is again lowest in the case of dynamic borrowing.



**TABLE 3** Overview of different treatment effect assumptions. The priors  $\gamma_{MonoA}$ ,  $\gamma_{MonoB}$  and  $\gamma_{Combo}$  as described in section A.3 are all pointwise with a support of 1,2 or 3 different points (each with probability “p”) and result in effective response rates  $\pi_{SoC}$ ,  $\pi_{MonoA}$ ,  $\pi_{MonoB}$ , and  $\pi_{Combo}$ . Only results of setting 1 are shown in the main text, the rest in the supplements

Setting	$\pi_{SoC}$	$\pi_{MonoA}$ ( $\gamma_{MonoA}$ )	$\pi_{MonoB}$ ( $\gamma_{MonoB}$ )	$\pi_{Combo}$ ( $\gamma_{Combo}$ )	Description
1	0.10	0.20 (2)	0.10 (1) with p = 0.5 0.20 (2) with p = 0.5	0.20 (1) if $\gamma_{MonoB} = 1$ 0.40 (1) if $\gamma_{MonoB} = 2$	Backbone monotherapy superior to SoC, add-on monotherapy has 50:50 chance to be superior to SoC; in case add-on monotherapy not superior to SoC, combination therapy as effective as backbone monotherapy, otherwise combination therapy significantly better than monotherapies
2	0.10	0.20 (2)	0.10 (1)	0.20 (1)	Backbone monotherapy superior to SoC, but add-on monotherapy not superior to SoC and combination therapy not better than backbone monotherapy
3	0.10	0.20 (2)	0.10 (1)	0.30 (1.5)	backbone monotherapy superior to SoC and combination therapy superior to backbone monotherapy, but add-on monotherapy not superior to SoC
4	0.10	0.20 (2)	0.10 (1)	0.40 (2)	Backbone monotherapy superior to SoC and combination therapy superior to backbone monotherapy (increased combination treatment effect compared to setting 4), but add-on monotherapy not superior to SoC
5	0.10	0.20 (2)	0.20 (2)	0.20 (0.5)	Both monotherapies are superior to SoC, but combination therapy is not better than monotherapies
6	0.10	0.20 (2)	0.20 (2)	0.30 (0.75)	Both monotherapies are superior to SoC and combination therapy is better than monotherapies
7	0.10	0.20 (2)	0.20 (2)	0.40 (1)	Both monotherapies are superior to SoC and combination therapy is superior to monotherapies (increased combination treatment effect compared to setting 7)
8	0.10	0.10 (1)	0.10 (1)	0.10 (1)	Global null hypothesis
9	0.20	0.20 (1)	0.20 (1)	0.20 (1)	Global null hypothesis with higher response rates
10	0.10	0.20 (2)	0.10 (1) with p = 0.5 0.20 (2) with p = 0.5	0.20* $\gamma_{MonoB}$ *0.5 (0.5) with p = $\frac{1}{3}$ 0.20* $\gamma_{MonoB}$ *1 (1) with p = $\frac{1}{3}$ 0.20* $\gamma_{MonoB}$ *1.5 (1.5) with p = $\frac{1}{3}$	Backbone monotherapy superior to SoC, add-on monotherapy has 50:50 chance to be superior to SoC; combination therapy interaction effect can either be antagonistic/non-existent, additive or synergistic (with equal probabilities)
11	0.10 + 0.03* (c-1)	0.10 + 0.03* (c-1) (1)	0.10 + 0.03* (c-1) (1)	0.10 + 0.03*(c-1) (1)	Time-trend null scenario; every new cohort (first cohort c = 1, second cohort c = 2, ...) will have SoC response rate that is by 3%-points higher than that of the previous cohort
12	0.10 + 0.03* (c-1)	0.20 + 0.03* (c-1) (2)	0.20 + 0.03* (c-1) (2)	0.40 + 0.03*(c-1) (1)	Time-trend scenario, whereby monotherapies superior to SoC and combination therapy superior to monotherapies; every new cohort (first cohort c = 1, second cohort c = 2, ...) will have SoC response rate that is by 3%-points higher than that of the previous cohort
13	0.20	0.30 (1.5)	0.30 (1.5)	0.40 ( $\frac{8}{9}$ )	Analogous to setting 7, but SoC response rate is 20%
14	0.20	0.30 (1.5)	0.30 (1.5)	0.50 ( $\frac{10}{9}$ )	Analogous to setting 8, but SoC response rate is 20%



**FIGURE 3** Impact of the data sharing (linetype and point shape) and maximum number of cohorts per platform (x-axis) on the per-cohort and per-platform type 1 error (Figure 3A) and power (Figure 3B) in treatment efficacy setting 1. Please note that different scaling of the y-axis is used in the two subfigures. With increasing number of cohorts in the platform, the chance to make at least one correct positive or negative decision increases. When data is shared, the per-cohort power increases, while it stays constant when no data sharing is planned

In Figure 4, we investigate the impact of optimistic/pessimistic assumptions regarding the expected platform size, cohort inclusion rate and final cohort sample size on per-cohort, and per-platform power. We observe that the PCP is independent of the assumptions regarding the maximum number of included cohorts and the cohort inclusion rate. When the sample size is increased, both PCP and disjunctive power increase, however the increase is faster when using more optimistic assumptions regarding the expected platform size and cohort inclusion rate, as firstly more data can be shared and secondly more cohorts will be included on average, thereby increasing the chance for any true positive decision.


Next, we removed the 50:50 chance for new cohorts entering the platform trial to have an efficacious add-on monotherapy and increased this probability to 100% (i.e., we are using treatment efficacy setting 7 instead of setting 1). We also wanted to investigate the impact of an increase in SoC response rate, so we increased the SoC response rate from 10% to 20% but kept the incremental increases in response rate the same (i.e., 10% points increased for the monotherapies and 30% points increased for the combination therapy; setting 14). Results for the power are presented in Figure 5. While all the relative influences of level of data sharing, use of decision rules, and so on, appear to be unchanged, we observed consistently slightly lower power for increased SoC response rate, which could be due to increased variance in the observed response rates. Another phenomenon that is particularly pertinent in this figure: While the PCP is always lowest when sharing no data, the disjunctive power is only the lowest when sharing no data when using small sample sizes. With increasing sample size the disjunctive power is greatest when sharing no data. While this might seem not intuitive at first, we believe it is due to random highs in the SoC (leading to unsuccessful comparisons of monotherapies vs. SoC) and backbone monotherapy (leading to unsuccessful comparisons of the combination therapy vs. backbone monotherapy) arms. When these occur, sharing this data across cohorts might negatively impact all other truly efficacious cohorts leading to simultaneously only false negative decisions. As a result, the disjunctive power drops. This cannot happen when no data is shared.

### 3.2.1 | Impact of modified decision rules


We further investigated the impact of parameters required for the Bayesian GO decision rules on the error rates. Every individual comparison includes a superiority margin  $\delta$  and a required confidence  $\gamma$  (e.g., posterior  $(P(\pi_{Comb} > \pi_{MonoA} + \delta | Data) > \gamma))$ . In the decision rules in section 2.2 we fixed  $\delta = 0$  and  $\gamma = 0.90$ . We now varied  $\gamma$  from 0.65 to 0.95 and  $\delta$  from 0 to 0.2. The impact on the operating characteristics is presented in Figure 6. Figure 6A reveals that both when sharing no or all data, the FWER can increase beyond 0.05 and even far beyond 0.10. Similarly, Figure 6B reveals that when choosing more lenient decision rules and increasing the data sharing, a PCP of close to 1 is attainable. Such contour plots will help to fine tune the Bayesian decision rules in order to achieve the desired operating characteristics. For  $\gamma$ , one would usually expect values equal to or larger than 80%. The choice of  $\delta$  also relies on clinical judgment. Generally, as a result of the combined decision rule, operating characteristics are rather conservative with respect to type 1 error rates.

### 3.2.2 | Summary of main results

Planning a platform trial is much more complex than planning a traditional full development trial. Firstly, both power and type 1 error might have to be controlled on the platform or cohort level. A platform trial that replaces sponsors' individual trials with the aim of advancing as many compounds as possible will try to control error rates on the cohort level. However, a platform trial with the goal of making sure that at least one efficacious compound is advanced would rather control error rates on the platform level. As seen in Figure 4, the per-cohort power depends on the assumptions regarding the cohort inclusion rate and number of cohorts that will enter over time. In case there is substantial uncertainty regarding how fast and how many new cohorts will enter the platform trial over time when



**FIGURE 4** Per-cohort and per-platform power with respect to data sharing (linetype and point shape), assumptions regarding the maximum number of cohorts (rows), cohort inclusion rate (columns) and final cohort sample size (x-axis) in treatment efficacy setting 1. Generally, both types of power increase with increasing final cohort sample size. In case of no data sharing, the per-cohort power is independent of assumptions regarding the expected platform size and cohort inclusion rate

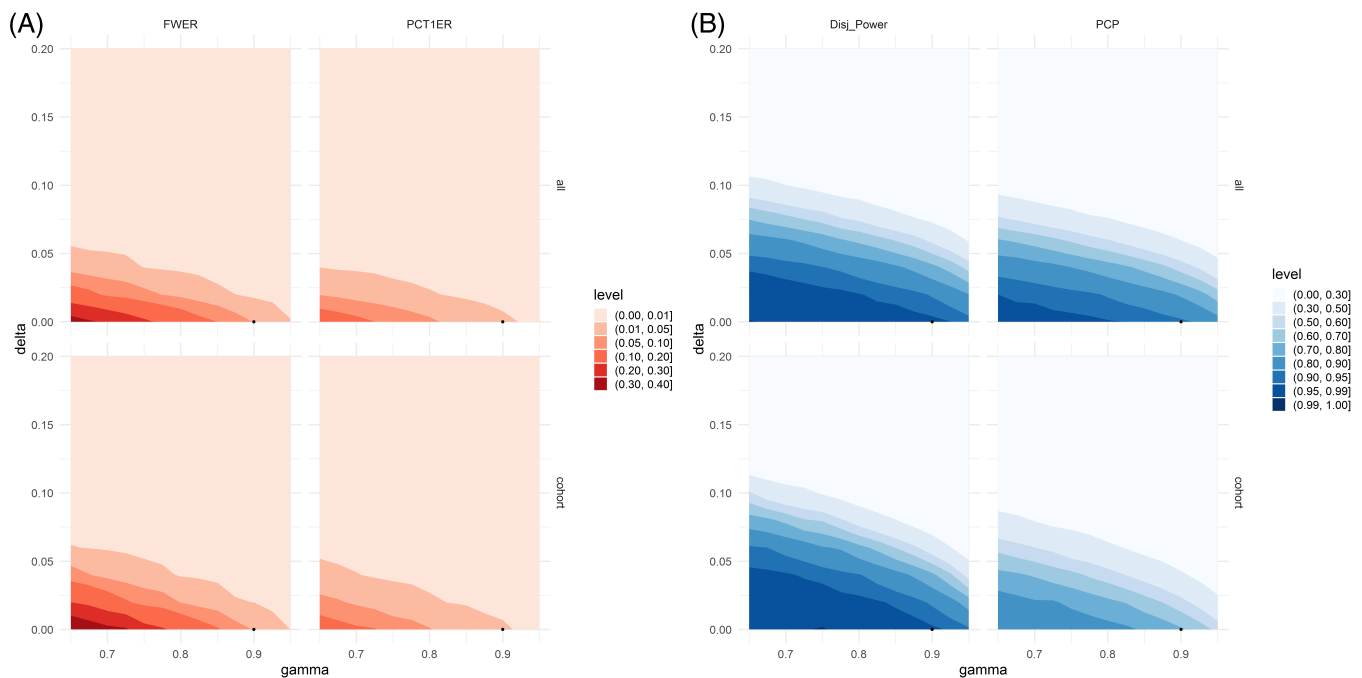


**FIGURE 5** Impact of SoC response rate (columns; treatment efficacy settings 7 vs. 14), final cohort sample size (x-axis) and data sharing (linetypes and point shapes) on per-cohort and per-platform power. We observed consistently lower power for increased SoC response rate. We also observed that while for lower sample sizes the per-platform power is increased with increasing amount of data sharing, this is not true anymore for larger sample sizes

determining the sample size, conservative lower bounds for the per-cohort power can be used by assuming no data sharing will take place. In this case, further simulations revealed that a final cohort sample size of around 600 would be needed to achieve a PCP of 0.8. In order to achieve a disjunctive power of 0.8, a final cohort sample size in the range of 300–500 is needed (depending on the assumptions; for few cohorts and a lower probability of cohort inclusion roughly 500 and for more cohorts and a higher probability of inclusion slightly below 300). In the fortunate event that the platform will attract many cohorts, the power for future cohorts would only increase with data sharing while guaranteeing a certain power for the first cohorts. If the priority is PCP and we are extremely confident about new cohorts entering the platform over time, by simply pooling all data we might achieve a PCP of 0.8 with a final cohort sample size of around 340 and a disjunctive power of 0.8 with a final cohort sample size of around 220. However, the above results are only true for the chosen treatment efficacy assumptions. When we are more optimistic about the treatment efficacy, the best choice for maximizing disjunctive power might be to not share any data at all (as seen in Figure 5, where in the left panel a disjunctive power of 0.8 is achieved with a final cohort sample size below 200 and no data sharing).

## 4 | DISCUSSION

To our knowledge, we conducted the first cohort platform simulation study to evaluate combination therapies with an extensive range of simulation parameters that reflect the potential complexity and a priori unknown trajectory of platform trials. The design under investigation is an open-entry, cohort platform study with a binary endpoint evaluating the efficacy of a number of two-compound combination therapies with a common backbone monotherapy compared to the respective monotherapies with putative individual efficacy over SoC. For example, the SoC arm could include a placebo add-on to achieve blinding. A range of treatment effect scenarios, as well as types of data sharing were



**FIGURE 6** Impact of the Bayesian GO decision rules on (A) per-cohort and per-platform type 1 error rates and (B) per-cohort and per-platform power in treatment efficacy setting 1. The black dot corresponds to the decision rules chosen in section 2.2. For every error rate, we set the data sharing (rows) to either full (“all”) or none (“cohort”). The x-axis shows the required confidence (“gamma”) and the y-axis the required superiority (“delta”) used in the Bayesian decision making in section 2.2. It is apparent that by choice of  $\delta$  and  $\gamma$ , a wide range of different error rates can be obtained

investigated, from no sharing to full sharing. In our simulations, no a priori knowledge is required regarding the possible trajectories of the platform trials, as in every simulation, different trajectories are dynamically simulated based on randomly generated events such as patients, outcomes, cohort inclusions, and so on. This means in one simulation, the platform trial might be stopped after including only one cohort and in another simulation run, it might stop after evaluating 7 cohorts. This is also a major difference compared to simulation programs for multi-arm multi-stage studies, which usually assume a fixed number of different arms or compute only arm-wise operating characteristics. Within each simulation in our program, periods of overlap between cohorts are recorded. We believe that especially this dynamic simulation is one of the key strengths of our simulation study, as it produces different realistic trajectories and computes operating characteristics across all such possible trajectories. We developed an R package and Shiny app alongside that facilitate easy usage and reproducibility of all simulation code, while the Shiny app for result exploration facilitates reproduction of our extensive simulation results, which we could not all discuss in this paper. Development of software for platform trial simulation is not straightforward and time spent on programming increases rapidly when more realistic features such as staggered entry, adding and dropping of treatment arms or data sharing across treatments/cohorts are included.<sup>40</sup>

Our results indicate that—apart from the treatment efficacy assumptions—the method of data sharing, exact specification of decision rules and complexity of the platform trial (expressed as maximum number of cohorts and cohort inclusion rate) are most influential to the operating characteristics of the platform trial. As expected, in nearly all cases, pooling all data leads to the largest power and type 1 error rates, whereas no data sharing leads to the lowest power. Methods that do not pool all data, but either discount them or use only concurrently enrolled patients, might significantly increase the power while only marginally increasing or even decreasing the type 1 error. Furthermore, definition of error rates in the context of a cohort platform trial with a combined decision rule per cohort is not straightforward.<sup>36,37,41–44</sup> As an example, whether or not to include simulated platforms that—as a result of random sampling of treatment efficacy from a specified prior distribution—do not contain truly efficacious cohorts in the calculation of the disjunctive power can lead to discrepancies of up to 15% points in our simulations. In terms of type 1 error rates, we defined a type 1 error via a target product profile and focused on interpreting the per-cohort type 1 error and the family-wise error rate in this paper. Other authors have recently suggested that FDR control might yield better

properties than family-wise error rate control in the multi-arm context.<sup>45</sup> The control of error rates such as the FDR might be appropriate for exploratory platform trials, which are less of a regulatory concern. Furthermore, in our investigations we treat configurations of “partial” superiority (e.g., the case where the combination is declared superior to the two monotherapies, but one or both of these are not declared superior to SoC) as non-rejections of the null hypothesis. Reversely, when for any of the four pair-wise comparisons conducted the null hypothesis holds, we consider any successful decision on the cohort level a false positive decision. For some configurations of null and alternative hypotheses within a cohort (e.g., when the combination therapy is superior to both monotherapies and SoC, but one of the components is not superior to SoC), the per-cohort type 1 error and the family-wise error rate are much higher than in the global null settings. For these configurations, partial power concepts might be more appropriate, or more relaxed testing strategies or thresholds than the one implemented could be considered, for example, if superiority of the backbone over SoC does not have to be shown within the platform trial. In this case, three instead of four pairwise comparisons would be required. Also it might make sense to use more relaxed thresholds for the comparisons of monotherapies versus SoC than for the comparisons of combination therapy versus the monotherapies. A further simplification would be to only test the combination therapy against the monotherapies (two pairwise comparisons) or against SoC (one pairwise comparison), in which case also the number of arms per cohort and allocation ratio would have to be adapted. For example in oncology, sometimes it is sufficient to demonstrate synergistic effects of the combination therapy, because monotherapies cannot separate from SoC. Such a combination therapy trial would be a success if combination therapy is superior to both monotherapies and SoC, even if the monotherapies do not separate from SoC.

As for any simulation program, especially in a highly complex and dynamic context such as platform trials, we had to make simplifications. The following restrictions must be acknowledged: (1) Patients are simulated in batches to achieve the targeted allocation rate. In more detail, for every time interval a batch of patients is simulated, all entering the study at the same time (imitating a block randomization). This led to potentially slightly different sample sizes at interim and final than planned. However, this is also not uncommon in reality. By simulating the outcomes this way, the overall run time could be significantly reduced, thereby allowing us to investigate more than ten thousand combinations of simulation parameters and simulating more than one hundred million individual platform trajectories. For more information on the exact patient sampling mechanism, see the R package vignette. (2) The dynamic borrowing approach we used is rather straightforward and heuristic. While based on a well-known and accepted method described by Schmidli et al.<sup>28</sup> we had to adapt the approach to be compatible with our cohort platform design and simulation structure, leading to some heuristic adaptations for the sake of computational efficiency. (3) Furthermore, in the simulations individual patient demographic or baseline characteristics were not considered. We assumed that patients were always drawn from the same population and different cohorts shared common inclusion/exclusion criteria. All newly simulated patients under such assumptions were eligible for randomization to any cohort ongoing at the time the patient was simulated, which enabled naive borrowing of data across cohorts during the analyses. In practice, separate cohorts could have some different inclusion / exclusion criteria due to scientific (e.g., treatments of different modes of action) or operational considerations. Nevertheless, the schematic approaches in this simulation study for assessing trial operating characteristics could remain instrumental. Otherwise, partial instead of complete data sharing could be planned for when designing a platform study. In such practical scenarios, prior epidemiological and clinical knowledge on the disease of interest also plays a critical role for appropriate data borrowing hence for the reliable assessment of power and error rates. (4) As we would expect in real life, our simulations show variations in the final sample size. Especially when performing more data sharing, the originally planned final sample size tends to be exceeded. Further simplifications included not simulating variations in recruitment speed due to availability of centers, external and internal events such as approval of competitor drugs or discontinuation of drug development programs. While the software does facilitate stopping treatments due to safety events, we did not use this parameter in the presented simulation study. As the best possible compromise under uncertainty, we aimed to achieve a balanced randomization for every comparison in case of either no data sharing or sharing only concurrent data and refrained from investigating response-adaptive randomization, which might increase the efficiency of the platform trial.<sup>13,46</sup> We only investigated platform trials with a maximum number of cohorts between 3 and 7. We believe this is sufficient in the initial planning step. If less cohorts are expected, a platform trial is not warranted. If more cohorts are expected to be included over time, simulations would need to be re-run.

To evaluate combination therapies the proposed platform design will be an efficient way to evaluate potential drugs and the resulting therapies. When exploring Bayesian decision rules, a key factor is fine-tuning of decision parameters (e.g.,  $\delta$  and  $\gamma$ ) in case operating characteristics should be controlled at a certain level. In the Bayesian decision rules, we tried to demonstrate superiority for the comparison of the monotherapies versus SoC. However, in some therapeutic areas, SoC could be an active drug with a different mechanism of action. In such instances, it could be sufficient to

demonstrate non-inferiority for some of the pairwise comparisons using pre-specified non-inferiority margins  $\delta < 0$ . Furthermore, if drugs have more than one dose level, the platform trial can be extended to allow for more dose levels per cohort. This may result in even more complex models and decision rules, depending on which data sharing models are to be used when calculating the posteriors. Considering dose–response relationships in combination trials will be a part of future research, as well as investigation of non-inferiority decision rules and impact of different allocation ratios on the operating characteristics.

## ACKNOWLEDGMENTS

The authors thank Constantin Kumaus for his support in programming the R Shiny App for online result exploration.<sup>39</sup>

## CONFLICT OF INTEREST

The research of ELM was funded by Novartis on the university and not an individual level. PM, CD-B, EG, YL are employees of Novartis Pharmaceuticals Corporation. FK reports grants on a university level from Novartis Pharma AG during the conduct of the study and from Merck KGaA outside the submitted work.

## DATA AVAILABILITY STATEMENT

For a complete overview of all simulation results, we developed a Shiny App facilitating self-exploration of all of our simulation results. The purpose of this R Shiny app is to quickly inspect and visualize all simulation results that were computed for this paper. We uploaded the R Shiny app, alongside all of our simulation results used in this paper to our server [39].

## ORCID

Elias Laurin Meyer  <https://orcid.org/0000-0001-5398-6334>

Peter Mesenbrink  <https://orcid.org/0000-0002-5624-6726>

Cornelia Dunger-Baldauf  <https://orcid.org/0000-0001-5917-7052>

Ekkehard Glimm  <https://orcid.org/0000-0003-3624-961X>

Yuhan Li  <https://orcid.org/0000-0003-4271-1540>

Franz König  <https://orcid.org/0000-0002-6893-3304>

## REFERENCES

1. Kunz CU, Jörgens S, Bretz F, et al. Clinical trials impacted by the covid-19 pandemic: adaptive designs to the rescue. *Stat Biopharmaceut Res*. 2020;12(4):461-477. doi:10.1080/19466315.2020.1799857
2. Stallard N, Hampson L, Benda N, et al. Efficient adaptive designs for clinical trials of interventions for covid-19. *Stat Biopharmaceut Res*. 2020;0:1-15. doi:10.1080/19466315.2020.1790415
3. Dodd LE, Follmann D, Wang J, et al. Endpoints for randomized controlled clinical trials for covid-19 treatments. *Clin Trials*. 2020;17(5):472-482.
4. Horby PW, Mafham M, Bell JL, et al. Lopinavir–ritonavir in patients admitted to hospital with covid-19 (recovery): a randomised, controlled, open-label, platform trial. *Lancet*. 2020;396(10259):1345-1352.
5. Angus DC, Derde L, Al-Beidh F, et al. Effect of hydrocortisone on mortality and organ support in patients with severe covid-19: the remap-cap covid-19 corticosteroid domain randomized clinical trial. *JAMA*. 2020;324(13):1317-1329.
6. Macleod J, Norrie J. Principle: a community-based covid-19 platform trial. *Lancet Respir Med*. 2021;9(9):943-945.
7. Woodcock J, LaVange LM. Master protocols to study multiple therapies, multiple diseases, or both. *N Engl J Med*. 2017;377(1):62-70. doi:10.1056/NEJMra1510062
8. Angus DC, Alexander BM, Berry S, et al. Adaptive platform trials: definition, design, conduct and reporting considerations. *Nat Rev Drug Discov*. 2019;18(10):797-808. doi:10.1038/s41573-019-0034-3
9. Park JJ, Siden E, Zoratti MJ, et al. Systematic review of basket trials, umbrella trials, and platform trials: a landscape analysis of master protocols. *Trials*. 2019;20(1):1-10. doi:10.1186/s13063-019-3664-1
10. Park JJ, Harari O, Dron L, Lester RT, Thorlund K, Mills EJ. An overview of platform trials with a checklist for clinical readers. *J Clin Epidemiol*. 2020;125:1-8.
11. Meyer EL, Mesenbrink P, Dunger-Baldauf C, et al. The evolution of master protocol clinical trial designs: a systematic literature review. *Clin Ther*. 2020;42(7):1330-1360. doi:10.1016/j.clinthera.2020.05.010
12. Israel E, Denlinger LC, Bacharier LB, et al. Precise: precision medicine in severe asthma: an adaptive platform trial with biomarker ascertainment. *J Allergy Clin Immunol*. 2021;147(5):1594-1601.
13. Saville BR, Berry SM. Efficiencies of platform clinical trials: a vision of the future. *Clin Trials*. 2016;13(3):358-366. doi:10.1177/1740774515626362

14. Brueckner M, Titman A, Jaki T, Rojek A, Horby P. Performance of different clinical trial designs to evaluate treatments during an epidemic. *PLOS One*. 2018;13(9):e0203387. doi:10.1371/journal.pone.0203387
15. Yuan Y, Guo B, Munsell M, Lu K, Jazaeri A. MIDAS: a practical Bayesian design for platform trials with molecularly targeted agents. *Stat Med*. 2016;35(22):3892-3906. doi:10.1002/sim.6971
16. Hobbs BP, Chen N, Lee JJ. Controlled multi-arm platform design using predictive probability. *Stat Methods Med Res*. 2018;27(1):65-78. doi:10.1177/0962280215620696
17. Tang R, Shen J, Yuan Y. CompAS: a Bayesian drug combination platform trial design with adaptive shrinkage. *Stat Med*. 2018;38(7):1120-1134. doi:10.1002/sim.8026
18. Ventz S, Alexander BM, Parmigiani G, Gelber RD, Trippa L. Designing clinical trials that accept new arms: an example in metastatic breast cancer. *J Clin Oncol*. 2017;35(27):3160-3168. doi:10.1200/JCO.2016.70.1169
19. FDA. Codevelopment of two or more new investigational drugs for use in combination. Guidance for Industry (June 2013). 2013. Accessed July 6, 2020. <https://www.fda.gov/media/80100/download>
20. EMA. Guideline on Clinical Development of Fixed Combination Medicinal Products (March 2017). 2017. Accessed July 6, 2020. [https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-clinical-development-fixed-combination-medicinal-products-revision-2\\_en.pdf](https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-clinical-development-fixed-combination-medicinal-products-revision-2_en.pdf)
21. Jiang L, Yan F, Thall PF, Huang X. Comparing bayesian early stopping boundaries for phase ii clinical trials. *Pharm Stat*. 2020;19(6):928-939.
22. Gallo P, Mao L, Shih VH. Alternative views on setting clinical trial futility criteria. *J Biopharm Stat*. 2014;24(5):976-993.
23. Stallard N, Todd S, Ryan EG, Gates S. Comparison of bayesian and frequentist group-sequential clinical trial designs. *BMC Med Res Methodol*. 2020;20(1):1-14.
24. Dawid AP. Selection paradoxes of bayesian inference. *Multivariate Anal Appl*. 1994;24:211-220.
25. Senn S. A note concerning a selection “paradox” of dawid’s. *Am Stat*. 2008;62(3):206-210.
26. Mandel M, Rinott Y. A selection bias conflict and frequentist versus bayesian viewpoints. *Am Stat*. 2009;63(3):211-217.
27. Viele K, Berry S, Neuenschwander B, et al. Use of historical control data for assessing treatment effects in clinical trials. *Pharm Stat*. 2014;13(1):41-54.
28. Schmidli H, Gsteiger S, Roychoudhury S, O’Hagan A, Spiegelhalter D, Neuenschwander B. Robust meta-analytic-predictive priors in clinical trials with historical control information. *Biometrics*. 2014;70(4):1023-1032.
29. Li W, Liu F, Snaveley D. Revisit of test-then-pool methods and some practical considerations. *Pharm Stat*. 2020;19(5):498-517.
30. Harun N, Liu C, Kim MO. Critical appraisal of bayesian dynamic borrowing from an imperfectly commensurate historical control. *Pharm Stat*. 2020;19(5):613-625.
31. Feißt M, Krisam J, Kieser M. Incorporating historical two-arm data in clinical trials with binary outcome: a practical approach. *Pharm Stat*. 2020;19(5):662-678.
32. Bofill Roig M, Krotka P & Glimm E. et al. On model-based time trend adjustments in platform trials with non-concurrent controls. arXiv preprint arXiv:2112.06574 (2021)
33. Bretz F, Koenig F, Brannath W, Glimm E, Posch M. Adaptive designs for confirmatory clinical trials. *Stat Med*. 2009;28(8):1181-1217.
34. Wason JMS, Stecher L, Mander AP. Correcting for multiple-testing in multi-arm trials: is it necessary and is it done? *Trials*. 2014;15(1):364. doi:10.1186/1745-6215-15-364
35. Howard DR, Brown JM, Todd S, Gregory WM. Recommendations on multiple testing adjustment in multi-arm trials with a shared control group. *Stat Methods Med Res*. 2018;27(5):1513-1530. doi:10.1177/0962280216664759
36. Stallard N, Todd S, Parashar D, Kimani PK, Renfro LA. On the need to adjust for multiplicity in confirmatory clinical trials with master protocols. *Ann Oncol*. 2019;30(4):506-509. doi:10.1093/annonc/mdz038
37. Collignon O, Gartner C, Haidich AB, et al. Current statistical considerations and regulatory perspectives on the planning of confirmatory basket, umbrella, and platform trials. *Clin Pharmacol Therapeut*. 2020;107(5):1059-1067.
38. Berry SM. Potential statistical issues between designers and regulators in confirmatory basket, umbrella, and platform trials. *Clin Pharmacol Therapeut*. 2020;108(3):444-446.
39. Meyer EL, Kumaus C. Shiny app for simulation result exploration and visualization. 2021. Accessed October, 12, 2021. <https://sny.cemsiis.meduniwien.ac.at/~zrx5rdf/oWer32/>
40. Meyer EL, Mesenbrink P, Mielke T, Parke T, Evans D, König F. Systematic review of available software for multi-arm multi-stage and platform clinical trial design. *Trials*. 2021;22(1):1-14.
41. Bai X, Deng Q, Liu D. Multiplicity issues for platform trials with a shared control arm. *J Biopharm Stat*. 2020;30(6):1077-1090.
42. Parker RA, Weir CJ. Non-adjustment for multiple testing in multi-arm trials of distinct treatments: rationale and justification. *Clin Trials*. 2020;17(5):562-566.
43. Bretz F, Koenig F. Commentary on parker and weir. *Clin Trials*. 2020;17(5):567-569.
44. Posch M, König F. Are p-values useful to judge the evidence against the null hypotheses in complex clinical trials? A comment on “the role of p-values in judging the strength of evidence and realistic replication expectations”. *Stat Biopharmaceut Res*. 2020;13:1-3.
45. Wason JM, Robertson DS. Controlling type i error rates in multi-arm clinical trials: a case for the false discovery rate. *Pharm Stat*. 2021;20(1):109-116.
46. Viele K, Saville BR, McGlothlin A, Broglio K. Comparison of response adaptive randomization features in multiarm clinical trials with control. *Pharm Stat*. 2020;19(5):602-612.
47. Kammer M. looplot: A package for creating nested loop plots; 2020. <https://github.com/matherealize/looplot>.
48. Rucker G, Schwarzer G. Presenting simulation results in a nested loop plot. *BMC Med Res Methodol*. 2014;14(1):129.



## SUPPORTING INFORMATION

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

**How to cite this article:** Meyer EL, Mesenbrink P, Dunger-Baldauf C, et al. Decision rules for identifying combination therapies in open-entry, randomized controlled platform trials. *Pharmaceutical Statistics*. 2022;21(3): 671-690. doi:10.1002/pst.2194

## APPENDIX A

### A.1 | Dynamic borrowing

Assume at any given time during the course of the platform trial we want to compute the posterior probability for either the backbone monotherapy or SoC response rate  $\theta_c$  in cohort  $c$  using a “dynamic borrowing” approach, in particular using a robust mixture prior (see Schmidli et al.<sup>28</sup>), in which we both use the cumulative observed data of cohort  $c$  (denoted by  $(n_c, k_c)$ ), as well as the cumulative pooled observed data from other cohorts in the platform trial (denoted by  $(n_c^p, k_c^p)$ ), where  $n$  denotes the sample size and  $k$  the number of responders. The data  $(n_c^p, k_c^p)$  might include pooled data from both earlier cohorts of the same platform trial, or concurrent data from concurrent cohorts of the same platform trial, but does not include data that has not yet been observed. For example when looking at Figure 2,  $(n_c, k_c)$  is the data from cohort 3 and  $(n_c^p, k_c^p)$  is the pooled data from cohorts 1 and 2 (for the backbone monotherapy and SoC arms, respectively). Ultimately our goal is to have a posterior for  $\theta_c$  that consists of two weighted Beta distributions, one informed by both  $(n_c, k_c)$  and  $(n_c^p, k_c^p)$  with weight  $w_1$  and the other informed only by  $(n_c, k_c)$  with weight  $w_2$  ( $w_1 + w_2 = 1$ ), in other terms

$$\pi(\theta_c) = w_1 * f(y, k_c^p + k_c + \alpha_c^p, n_c^p + n_c - k_c^p - k_c + \beta_c^p) + w_2 * f(y, k_c + \alpha_c, n_c - k_c + \beta_c), \quad (\text{A1})$$

where  $f(\cdot, \alpha, \beta)$  denotes the probability density function of a Beta distribution with parameters  $\alpha$  and  $\beta$  and  $\pi(\theta_c)$  denotes the density of the posterior distribution of  $\theta_c$ . The weights then represent the degree of “borrowing.” Furthermore, the a priori unknown degree of borrowing (i.e.,  $w_1$  and  $w_2$ ) should be based on  $(n_c, k_c)$  and  $(n_c^p, k_c^p)$ .

Our robust mixture prior for  $\theta_c$  conditional on this data is given by

$$\begin{aligned} \pi^{RMP}(\theta_c | n_c^p, k_c^p) &= w * f(\theta_c, k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c) + (1 - w) * f(\theta_c, \alpha_c, \beta_c) = \\ &= w \frac{\theta_c^{k_c^p + \alpha_c - 1} * (1 - \theta_c)^{n_c^p - k_c^p + \beta_c - 1}}{B(k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c)} + (1 - w) \frac{\theta_c^{\alpha_c - 1} * (1 - \theta_c)^{\beta_c - 1}}{B(\alpha_c, \beta_c)}, \end{aligned} \quad (\text{A2})$$

where  $B(\alpha, \beta)$  denotes the Beta function at  $(\alpha, \beta)$  and  $\alpha_c$  and  $\beta_c$  are the parameters reflecting the current information about  $\theta_c$  (both set to 0.5 through our simulations). Please note that this posterior is not iteratively derived, but rather “rebuilt” from scratch every time an analysis is conducted. Please further note that  $w_1$  and  $w_2$  (Equation (A1)) depend on  $w$  (Equation (A2)) and  $w$  needs to be chosen a priori. The posterior distribution of  $\theta_c$  is now derived as:

$$\begin{aligned} \pi(\theta_c) &\propto \pi^{RMP}(\theta_c | n_c^p, k_c^p) * \theta_c^{k_c} * (1 - \theta_c)^{n_c - k_c} = \\ &= w \frac{\theta_c^{k_c^p + k_c + \alpha_c - 1} * (1 - \theta_c)^{n_c^p + n_c - k_c^p - k_c + \beta_c - 1}}{B(k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c)} + (1 - w) \frac{\theta_c^{k_c + \alpha_c - 1} * (1 - \theta_c)^{n_c - k_c + \beta_c - 1}}{B(\alpha_c, \beta_c)}. \end{aligned} \quad (\text{A3})$$

In order for this to be a distribution, we need to make sure it integrates to 1. An option which captures the idea of dynamic borrowing in the sense that the more similar the posterior of  $\theta_c$  based on  $(n_c, k_c)$  is to the posterior of  $\theta_c$  based on  $(n_c^p, k_c^p)$ , the more borrowing should be done - is to set

$$w_1 = \frac{w \frac{B(k_c^p + k_c + \alpha_c, n_c^p + n_c - k_c^p - k_c + \beta_c)}{B(k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c)}}{w \frac{B(k_c^p + k_c + \alpha_c, n_c^p + n_c - k_c^p - k_c + \beta_c)}{B(k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c)} + (1-w) \frac{B(k_c + \alpha_c, n_c - k_c + \beta_c)}{B(\alpha_c, \beta_c)}} \quad (\text{A4})$$

and

$$w_2 = \frac{(1-w) \frac{B(k_c + \alpha_c, n_c - k_c + \beta_c)}{B(\alpha_c, \beta_c)}}{w \frac{B(k_c^p + k_c + \alpha_c, n_c^p + n_c - k_c^p - k_c + \beta_c)}{B(k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c)} + (1-w) \frac{B(k_c + \alpha_c, n_c - k_c + \beta_c)}{B(\alpha_c, \beta_c)}}. \quad (\text{A5})$$

This corresponds to the case in which we divide the unscaled posterior in Equation (A2) by the factor

$$w \frac{B(k_c^p + k_c + \alpha_c, n_c^p + n_c - k_c^p - k_c + \beta_c)}{B(k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c)} + (1-w) \frac{B(k_c + \alpha_c, n_c - k_c + \beta_c)}{B(\alpha_c, \beta_c)}. \quad (\text{A6})$$

The posterior distribution is now given by

$$\pi(\theta_c) = \frac{w \frac{\theta_c^{k_c^p + k_c + \alpha_c - 1} (1 - \theta_c)^{n_c^p + n_c - k_c^p - k_c + \beta_c - 1}}{B(k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c)} + (1-w) \frac{\theta_c^{k_c + \alpha_c - 1} (1 - \theta_c)^{n_c - k_c + \beta_c - 1}}{B(\alpha_c, \beta_c)}}{w \frac{B(k_c^p + k_c + \alpha_c, n_c^p + n_c - k_c^p - k_c + \beta_c)}{B(k_c^p + \alpha_c, n_c^p - k_c^p + \beta_c)} + (1-w) \frac{B(k_c + \alpha_c, n_c - k_c + \beta_c)}{B(\alpha_c, \beta_c)}}. \quad (\text{A7})$$

Since  $w_1 + w_2 = 1$ , we can be sure that this yields a proper distribution. Choosing the weights  $w_1$  and  $w_2$  in this way now leads to the borrowing being dynamic. In Figure 7 we investigate the impact of the sample sizes and observed response rates in the cumulative observed cohort and cumulative pooled observed data from other cohorts in the platform trial, as well as the prior  $w$  on the weight  $w_1$ .

Finally, the parameters of the Beta posterior are derived as


$$\alpha_{\text{eff}} = w_1 * (k_c^p + k_c + \alpha_c) + w_2 * (k_c + \alpha_c) \quad (\text{A8})$$

$$\beta_{\text{eff}} = w_1 * ((n_c + n_c^p) - (k_c^p + k_c) + \alpha_c) + w_2 * (n_c - k_c + \alpha_c) \quad (\text{A9})$$

Despite not being investigated in this manuscript, the simulation software also features frequentist decision rules based on  $2 \times 2$  contingency tables and therefore expects the effective sample size and number of responders to be integers. We therefore simplify and assume the effective sample size to be  $\lfloor \alpha_{\text{eff}} + \beta_{\text{eff}} \rfloor$  and the effective number of responders to be  $\lfloor \alpha_{\text{eff}} \rfloor$ . The whole procedure is applied for backbone monotherapy and SoC separately.

## A.2 | Different target product profiles

All considered pair-wise decision rules are of the form “GO,” if  $P(\pi_y > \pi_x + \delta | \text{Data}) > \gamma$  (see section 2.2). In section 2.4, we assumed that the target product profile would always require the response rate of drug  $y$  to be greater than the response rate of drug  $x$  by any margin to consider the alternative hypothesis to hold. This can be extended to allow the target product profile to specify any required margin  $\zeta$ , such that if and only if the response rate of drug  $y$  is truly better than the response rate of drug  $x$  by a margin of at least  $\zeta$  do we consider the alternative hypothesis to hold. For every pair-wise comparison, this margin  $\zeta$  can, but does not have to, coincide with the  $\delta$  used. This is visualized for a single pair-wise comparison in Figure 8. As a reminder, for all of the pair-wise comparisons in our simulations, we set  $\zeta = 0$ . In the **CohortPlat** package, a different margin  $\zeta$  can be chosen both for the comparison of combination against the monotherapies and the comparison of the monotherapies against SoC.



**FIGURE 7** Impact of the sample sizes and observed response rates in the cumulative observed cohort and cumulative pooled observed data from other cohorts in the platform trial ( $n_c, \pi_c = \frac{k_c}{n_c}$  and  $n_p, \pi_p = \frac{k_p}{n_p}$  respectively), as well as the prior  $w$  on the weight  $w_1$ . Please note in the figure the label “pi\_c” is used for  $\pi_c$ , “pi\_p” is used for  $\pi_p$ , “nc” is used for  $n_c$  and “np” is used for  $n_p$ . When the observed response rates in the two groups are the same, the amount of data sharing increases with  $w$  and increasing sample size in both groups. When the observed response rates are different, the amount of data sharing decreases with decreasing  $w$  and increasing sample size in both groups. When the observed response rates are very different, the amount of data sharing is nearly 0, even when the sample sizes are small and the prior  $w$  was chosen to be 0.9. When, for example, the prior  $w$  is set to 0.1, even in case of equality of  $\pi_c$  and  $\pi_p$ , the actual borrowing  $w_1$  does not increase beyond approximately 0.7 for the chosen sample sizes. For all results, the Beta prior parameters were set to 0.5. Visualization is based on the **loopt** package,<sup>47</sup> which implements the visualization presented by Rucker and Schwarzer.<sup>48</sup> Please note that, for example, in the plot in the bottom left corner, most lines are overlapping

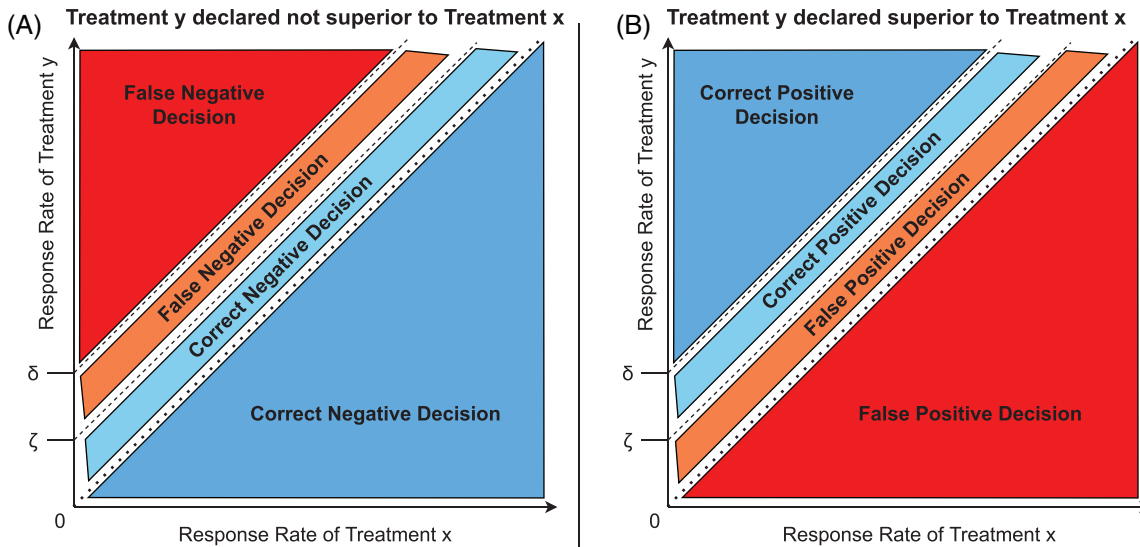
**A.3 | Treatment efficacy scenarios**

In the main text, we showed nearly exclusively selected results of one set of assumptions regarding the treatment effect of the monotherapies and the combination treatment (setting 1). In general, we only investigated treatment effect assumptions based on risk-ratios, whereby we randomly and separately draw the risk-ratio for each of the monotherapies with respect to the SoC treatment. For the combination treatment, we randomly draw from a range of interaction effects, which could result in additive, synergistic or antagonistic effects of a specified magnitude. Some scenarios might be more realistic for a given drug development program than others, however we felt that the broad range of scenarios will allow to investigate the impact and interaction of the various simulation parameters and assumptions on the operating characteristics. Let  $\pi_x$  denote the probability of a patient on therapy  $x$  to have a successful treatment outcome (binary), that is, the response-rate, and  $T_x$  denote a discrete random variable.

In detail, every time a new cohort enters the platform, we firstly fix the SoC response-rate:

$$\pi_{SoC} \in [0, 1]$$

Then we assign the treatment effect in terms of risk-ratios for the backbone monotherapy (monotherapy A), which is the same across all cohorts:



**FIGURE 8** After a pair-wise comparison of treatment x against treatment y, whereby the decision rules required the response rate of treatment y ( $\pi_y$ ) to be superior to the response rate of treatment x ( $\pi_x$ ) by a margin of  $\delta$  (i.e.,  $\pi_y \geq \pi_x + \delta$ ), treatment y was either declared superior (subfigure B) or not superior (subfigure A) to treatment x. A target product profile was defined for treatment y, whereby the aim was for the response rate of treatment y to be superior to the response rate of treatment x by a margin of  $\zeta$  (i.e.,  $\pi_y \geq \pi_x + \zeta$ ). For simplicity, we assume  $\delta \geq \zeta$ , although this approach also works when  $\delta < \zeta$ . When in truth  $\pi_y < \pi_x + \zeta$ , the decision is either a correct negative decision (subfigure A) or a false positive decision (subfigure B). When in truth  $\pi_y \geq \pi_x + \zeta$ , the decision is either a false negative decision (subfigure A) or a correct positive decision (subfigure B). Please note that in the usual definition of a type 1 error,  $\zeta = 0$

$$\pi_{MonoA} = \pi_{SoC} * \gamma_{MonoA} \cdot \gamma_{MonoA} \sim T_{MonoA}$$

Then we randomly draw the treatment effect in terms of risk-ratios for the add-on monotherapy (monotherapy B):

$$\pi_{MonoB} = \pi_{SoC} * \gamma_{MonoB} \cdot \gamma_{MonoB} \sim T_{MonoB}$$

Finally, after knowing the treatment effects of both monotherapies, we randomly draw an interaction effect for the combination treatment:

$$\pi_{Combo} = \pi_{SoC} * (\gamma_{MonoA} * \gamma_{MonoB}) * \gamma_{Combo} \cdot \gamma_{Combo} \sim T_{Combo}$$

Depending on the scenario, the distribution functions can have all the probability mass on one value, that is, the assignment of treatment effects and risk-ratios is not necessarily random. Please further note that while the treatment effects were specified in terms of risks and risk-ratios, the Bayesian decision rules were specified in terms of response rates. Two settings characterize global null hypotheses, six settings characterize an efficacious backbone monotherapy with varying degrees of add-on mono and combination therapy efficacy, two settings characterize an efficacious backbone with varying degrees of random add-on mono and combination therapy efficacy, two settings characterize either the global null hypothesis or efficacious mono and combination therapies, but with an underlying time-trend, and two settings were run as sensitivity analyses with increased standard-of-care response rates. The different treatment efficacy settings are summarized in Table 3.