Maximizing the Efficiency of Proof of Concept Studies and of Arrays of Proof of Concept Studies for Multiple Drugs or Indications

Robert A. Beckman, MD[1], Linchen He, MS[2,3],
[1]Departments of Oncology and of Biostatistics, Bioinformatics, and Biomathematics, Lombardi Comprehensive Cancer Center and Innovation Center for Biomedical Informatics, Georgetown University Medical Center, 2115 Wisconsin Avenue, Suite 110, Washington DC 20007
[2]Department of Biostatistics, Bioinformatics, and Biomathematics, Innovation Center for Biomedical Informatics, Georgetown University Medical Center, 2115 Wisconsin Avenue, Suite 110, Washington, DC 20007
[3]Current address: Department of Population Health, New York University, 650 1st Avenue, New York, New York 10016

Abstract
Phase II proof-of-concept (POC) trials determine which therapeutic hypotheses will undergo definitive Phase III testing. The number of possible POC hypotheses likely far exceeds available public or private resources. We propose a design strategy for maximizing efficiency of POC trials that obtains the greatest knowledge with the minimum patient exposure. We compare efficiency using the benefit-cost ratio, defined to be the risk-adjusted number of truly active drugs correctly identified for Phase III development divided by the risk-adjusted total sample size in Phase II and III development. It is most cost-effective to conduct smaller POC trials which are powered at 80% on an effect size 50% larger than that of minimal clinical interest, allowing more possible POC hypotheses to be investigated under a finite POC budget constraint. We also consider parallel arrays of POC trials with multiple indications or drugs, and sequential two-stage POC trial arrays where all drugs get an initial allotment of POC trials and only those which achieve a POC get further investment. These strategies can improve the output of successful drugs by up to 30% at a constant budget.

Key Words: utility functions, development efficiency, proof of concept study, optimization

1. Introduction

The cost of drug development is extremely high, on the order of $1 billion per successful new drug/indication pair. Much of this cost is accounted for by late stage clinical development costs, including the costs of failed development

efforts. Thus an increase in efficiency of this process by even 10% would save on the order of $100 million per successful drug/indication pair.

Proof of Concept (PoC) trials determine which drug/indication pairs will go on to definitive Phase 3 testing. In oncology, which is the particular motivation for this work, there are usually a very large number of possible PoC trials, representing a combinatorial explosion of a large number of possible drugs, drug combinations, dosing schedules, tumor types and histological and molecular subtypes, and lines of therapy. While preclinical data and early clinical data give some information for prioritizing among PoC trials, this information is generally not as predictive as one would hope. Hence the number of potentially useful PoC trials generally exceeds available patients as well as available public or private funds.

Chen and Beckman [1-3] sought to optimize the efficiency of PoC trials and their associated Go/No-Go criteria for phase 3, studying single PoC trials as well as portfolios of PoC trials conducted in parallel subject to a budgetary constraint on the PoC budget. They developed strategies which could increase the efficiency of late stage drug development by 10-30%.

In this paper, we will first summarize these previous results, outlining the principles of optimization, including the concept of Type III error, and presenting optimal designs for portfolios of parallel PoC studies. We will then discuss current applications of this work as well as questions and controversies which have arisen concerning this approach. We will then present preliminary findings of new research on 2 stage sequential arrays of PoC studies, wherein each stage may consist of several parallel PoC studies. Finally, we will suggest directions for future research.

2. Optimization of Parallel Arrays of PoC Studies

2.1 Principles of Optimization

This work was motivated by the observation by RAB that different pharmaceutical companies in which he had been employed had different approaches to PoC study design. Some companies conducted very small PoC trials in an effort to minimize cost per PoC trial. Others favored larger, more rigorous studies with strict control of Type I and II error. Development teams created clinical development plans to mimic the clear and subjective preferences of their management.

Instead, RAB proposed that development strategy should be informed by **quantitative, objective optimization of utility.** The proposed utility function for late development was a **benefit to cost ratio (BCR)**, also termed the **efficiency function**. Benefit was defined as the number of true positive PoC trials achieved by the development plan. Cost was defined in units of patients, and was the risk adjusted number of patient utilized in late development (ie Phase 2 and Phase 3), including Phase 3 studies performed in error due to false positive PoC studies. Mathematically, the efficiency function was relatively simple:

$$Efficiency = \frac{p\,(1-\beta)}{\lambda + p\,(1-\beta) + (1-p)\alpha}$$

(Equation 1),

where p is the probability that the true effect size associated with the drug is greater than or equal to the clinical effect size of interest, α and β are the Type I and II errors for the PoC trials, and λ is the ratio of the sample size of the PoC trial to the sample size of the definitive Phase III trial. Note λ is smaller for small PoC trials.

Optimization of this function was then performed for parallel arrays of PoC studies subject to a constrained PoC budget [1-3]. It was assumed that any PoC study which met the Go/No-Go bar would lead to a corresponding Phase 3 study. The Phase 3 studies did not have internal futility analyses, although this tactic could further increase the efficiency in practice.

At a constant PoC budget which is insufficient to execute all PoC trials of interest (the typical real world situation), there is a strategic trade-off between the number of PoC trials executed and their type I and type II errors. One may reduce the type I and type II error by performing larger PoC trials, but one will be able to test fewer worthy PoC hypotheses within the fixed budget.

In addition to the trade-off regarding PoC trial size, there is a trade-off concerning the Go-No Go criterion determining if development proceeds to Phase 3 after a PoC trial. A low bar for proceeding to Phase 3, which we termed the "No Drug Left Behind" strategy, maximizes the number of effective drugs produced, but it wastes considerable resources on Phase 3 trials that in the end prove to be based on false positive PoC studies. Importantly, when a large, expensive Phase 3 study fails, it is a highly visible outcome and may even have adverse career consequences for those who proposed it. In contrast, the more popular "Fail Early, Fail Cheap" strategy features a high bar to Phase 3. This avoids failed Phase 3 trials, but may lead to lost opportunities. However, a lost opportunity, either from a false negative PoC trial or from a PoC trial not performed due to budgetary limitations, is usually an invisible mistake.

Optimization attempts to find the best balance among the above strategic tradeoffs, and determines optimal Type I and II errors (and hence sample sizes) and Go-No Go criteria for proceeding to Phase 3, for each PoC trial in a portfolio of such trials executed in paralle1.

In performing the optimization, it became clear that in addition to the classical type I and II errors, a third phenomenon was important in determining the optima, and we termed it Type III error [4]. Type III error is the opportunity cost of not investigating worthy PoC hypotheses because of the limited PoC budget. If we choose to do larger PoC trials, we can reduce the Type I and II errors, but the Type III error will increase since we cannot do as many PoC trials. Conversely, by doing smaller PoC trials, we can do more of them, thus reducing the Type III error. However, smaller PoC trials will inevitably be associated with higher Type I and/or Type II errors.

The Type III error has a simple mathematical definition: it is the fraction of potentially true positive PoC outcomes "lost" because the corresponding PoC trials were not performed. It does not include the true positive PoC outcomes which were lost due to

false negatives (that is Type II error). Thus, the Type III error can be defined only at a constant Type II error:

$$Type\ III\ error = 1 - \frac{Expected\ number\ of\ true\ positive\ PoC\ outcomes\ |subset\ of\ PoC\ trials\ conducted}{Expected\ number\ of\ true\ positive\ PoCoutcomes|all\ PoC\ trials\ conducted}$$

(Equation 2),

where the quantity is defined only where the studies in the numerator and denominator are conducted at the same Type II error.

1.2 Optimal Designs for Parallel Portfolios of PoC studies

For the case of parallel PoC trials where all drug/indication pairs have equal values of p, a surprising result was found: while the optimal value of α was in the range of 0.05 -0.10, a traditional range for PoC trials, the optimal value of β was 0.4, twice the traditional value usually employed. This is because increasing β by this amount allowed twice as many PoC trials to be conducted for the same cost, substantially reducing the Type III error. The counterintuitive approach increased efficiency by 10-30% (depending on the effect sizes for Phase 2 and Phase 3 endpoints), thus potentially saving $100-300M per successful drug/indication pair.

The principle is illustrated graphically in Figure 1 below. In this figure, we assume we have two PoC opportunities each with p = 0.3, but only sufficient budget to execute one PoC trial at traditional α and β. In figure 1A, we present the scenario in which one large PoC trial is executed with β = 0.2, and in figure 1B the scenario in which two PoC trials are executed, each with β = 0.4.

The figures depict probabilistic branching pathways by which risk adjusted benefit and cost are computed. The first branch (on the left) is whether or not the drug is truly active. For p = 0.3, there is a 30% chance the drug is truly active, 70% that it is truly inactive. The next branch is the outcome of the PoC trial. If the drug is truly active, the possible outcomes are true positive, with probability 1 – β, and false negative, with probability β. If the drug is truly inactive, the possible outcomes are true negative, with probability 1 – α, and false positive, with probability α. The possible outcomes are given on the right. The chance of any outcome is the product of the arrows leading to it, from left to right. A true positive gives a benefit of 1, all other outcomes a benefit of zero. True and false negatives have a cost equivalent to the number of patients in the PoC trial (120 patients for the traditional PoC trial in figure 1a, 60 patients for the small PoC trials in figure 1b) , while true and false positives incur the cost of both the PoC trial and a subsequent Phase 3 trial (600 patients).

The risk adjusted benefit and cost are given at the bottom of the figures. With these parameter values, the single traditional PoC trial strategy requires 984

patients on average to identify and confirm a true positive drug/indication pair, whereas two smaller PoC trials require an average of 872 patients, or an 11% savings.

The benefit above is quite robust to the value of p. This is quite fortunate since p is difficult to objectively estimate. It is also robust to the relative costs of Phase 2 and Phase 3, and varies between 10% and 30%.

Figure 1A:

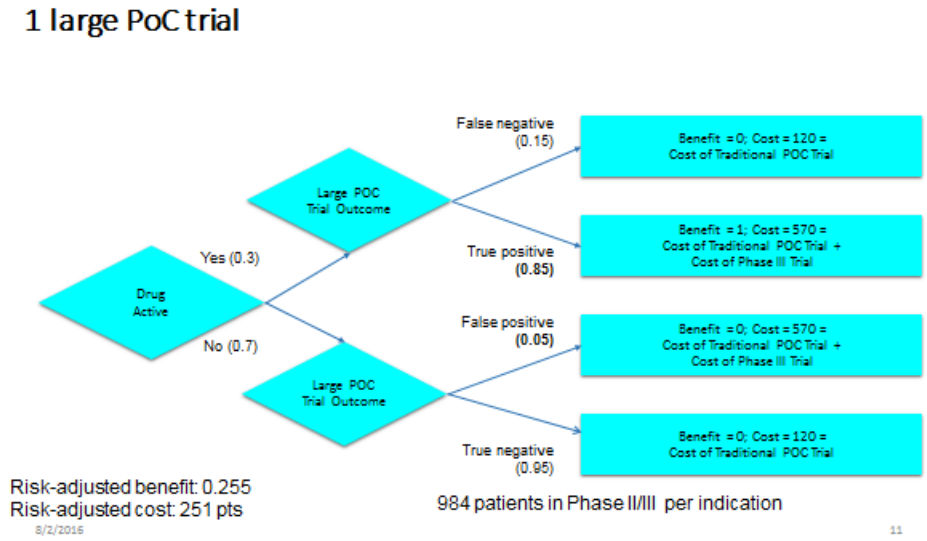

Figure 1B:

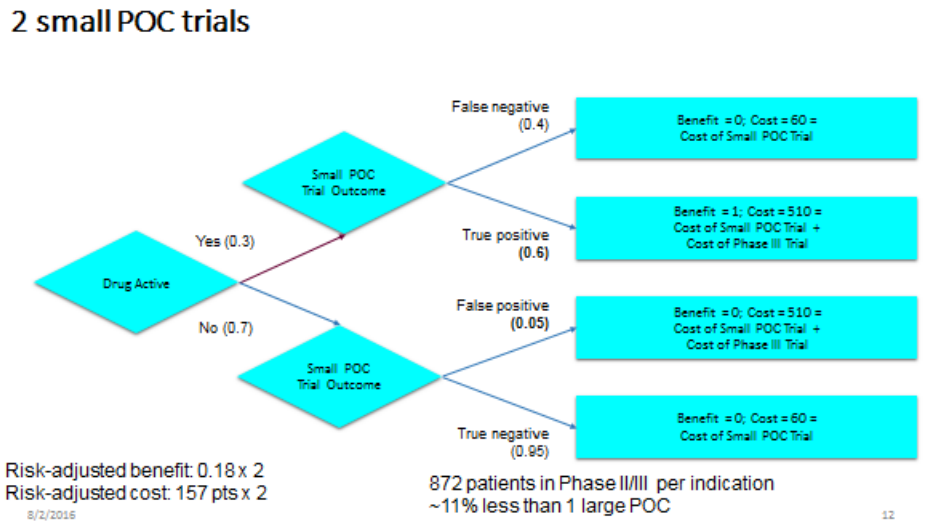

While the opportunity to increase the efficiency of drug development has been utilized extensively, we have occasionally encountered statisticians who cling to

the rule of thumb about which they were trained, $\beta = 0.2$. Interestingly, we may express our results in a different way: power the PoC study at 80%, but for an effect size 50% larger than the minimum effect size of clinical interest [5]. This formulation is exactly equivalent, since such a trial has 60% power for the minimum effect size of clinical interest, and an identical sample size and Go/No criteria to Phase 3 compared to the original recommendation. However, this formulation is more universally accepted since it does not overtly violate the rule of thumb $\beta = 0.2$. Moreover, powering on larger effect sizes is currently popular in oncology due to the optimism prevailing in the era of precision medicine. We discuss the perceptions of stakeholders here, because we wish to encourage a different mode of thinking. We wish to convey the point that the current traditional values of α and β are arbitrary, and are not derivable by logical argument. In designing drug development programs, the practitioner should make deliberate strategic choices of α and β, and objectively evaluate the performance of these choices relative to others according to quantitative criteria.

The results above assume the value of p is the same for all the drug/indication pairs, and that each drug/indication pair has the same utility if successful. However, straightforward extensions of the efficiency function are possible with different values of p for different indications, as well as weighting drug/indication pairs according to their relative value, expressed in units of patient benefit, population size, or economic impact, as desired [2,3]. These extensions allow this formalism to be used to optimize real portfolios of multiple indications for a drug, and also across multiple drugs in a pharmaceutical portfolio and their indications [6,7].

A simple example of portfolio optimization is given in Figure 2 below, representing a simple portfolio of two drug/indication pairs, one with a higher value of p than the other. Here the optimal strategy depends on the relative cost of the Phase 2 PoC trial to the cost of Phase 3. If the Phase 2 PoC trial cost is a relatively large fraction of the Phase 3 trial cost, the resources are allocated roughly proportionately to the value of p. However, as Phase 3 becomes relatively more expensive compared to a PoC trial, a breakpoint occurs, and the optimal strategy involves allocating all the resources to the drug with the superior value of p.

Figure 2:

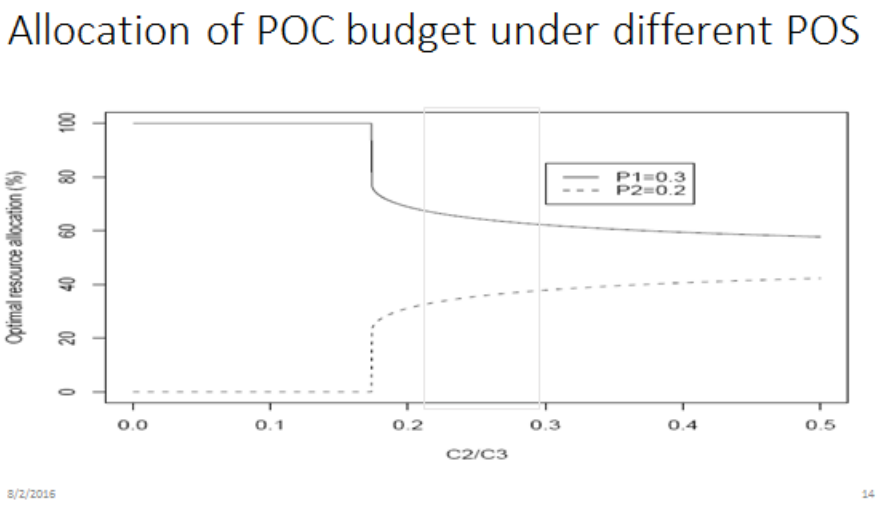

1.3 Applications of PoC Trial Optimization

The above approach has been used for sizing of PoC trials at Merck & Co, where it was developed. In order to gain universal acceptance at Institutional Review Boards, the formulation based on powering for a larger effect size and a traditional value of β was generally used.

In the era of precision medicine, phase 2 PoC trials must often be stratified based on the status of a putative predictive biomarker. Ideally, each stratum should be sufficiently powered to give an answer. Thus, the biomarker stratified PoC trial can be viewed as having two objectives: 1. Assess efficacy in the "biomarker positive" subset, and 2. Assess efficacy in the "biomarker negative" subset. Alternatively, one can express these two objectives as: 1. Assess the efficacy of the drug, and 2. Assess the predictive power of the biomarker. Either way, the predictive biomarker stratified PoC trial might be twice the size of an unstratified PoC trial. Fortunately, such a trial may be appropriately powered if each stratum is powered with $\beta = 0.4$ [4].

The benefit-cost ratio (BCR) approach has also been used to optimize the timing and futility threshold for futility analysis in Phase 3 trials, as well as to optimize Phase 2/3 seamless designs [6,8]. Interestingly, it has recently been discovered that the BCR should be considered when creating adaptive designs involving sample size readjustment. If consideration is not given to the effects of overpowering, sample size readjustment may actually reduce the overall efficiency of a clinical trial program [9].

1.4 Questions and Controversies Regarding PoC Trial Optimization

1.4.1 Why not use the absolute benefit, B-C, instead of the benefit/cost ratio B/C?

Indeed, exact optimization is in principle possible with the absolute benefit, whereas a benefit-cost ratio provides only practical guidance. However, in practice the quality of optimization of the absolute benefit is difficult to maintain, since B is difficult to estimate, due to uncertain and constantly changing conditions in the competitive environment, the regulatory environment, and the payer environment. The optimization of B – C is very sensitive to the value of B, especially when B is much larger than C. In contrast, errors in estimating B simply multiply the BCR by a constant factor but have minimal effect on the optimum.

Because the absolute benefit may continue to increase past the point of diminishing returns, it may lead to overspending if it is optimized for a subset of the portfolio. Only if absolute benefit is computed in a comprehensive manner across the entire portfolio will the absolute benefit be sensitive to spending past the point of diminishing returns within individual programs. The difficulty of simultaneously analyzing the entire portfolio, including the somewhat subjective and difficult estimates of B, should not be underestimated. In contrast, the BCR implicitly constrains each individual program not to spend past the point of diminishing returns.

1.4.2 Why not constrain the total Phase 2 and Phase 3 budget rather than only the Phase 2 PoC budget?

In our experience, while the PoC budget is fixed, diversified pharmaceutical companies generally find additional funds to support Phase 3 trials in the event of a positive PoC.

1.4.3 Some subsequent workers recommend rasing alpha rather than beta.

Subsequent workers using a similar approach have also recommended more, smaller PoC trials, but in contrast to our analysis they favor raising α rather than β [10]. We wish to first make the point that the most important message of this paper is that clinical development should be quantitately optimized according to objective utility functions. Specific recommendations may vary depending on goals and assumptions associated with the optimizarion.

Nonetheless, we find the recommendation in [10] intuitively implausible since a larger number of PoC trials at high α would lead to a flood of false positives and failed Phase 3 trials at great expense. Raising α can seem optimal if sunk preclinical and phase I costs are included in the optimization. However, most economists do not recommend including sunk costs in determining a strategy, since they are irreversible. Raising α may also seem optimal if the cost of Phase 3 trials decreases significantly in the future. The advent of precision medicine in oncology has led to some optimism that Phase 3 costs will decrease, and this may change the optimal strategy. Nonetheless, the current precision medicine paradigm does not fully account for molecular and phenotypic heterogeneity of

subclones within a single individual's cancer, nor for the potential of cancer cells to dynamically evolve therapeutic resistance [11]. As such, it is unclear whether the current approach to precision medicine will lead to overall reductions in Phase 3 costs, or merely be confined to a limited number of cases constituting low-hanging fruit.

3. Optimization of 2 Stage Arrays of PoC Trials

Many pharmaceutical companies do not fund their development teams for all recommended PoC trials in one stage. For example, out of N potentially valuable PoC trials for a given drug, budgetary constraints may lead to funding only a subset n of these trials. But all trials may not be funded for simultaneous parallel execution. Rather, the company may fund a subset of these trials, $n_1 < n$, in an initial stage. Only if at least one of the stage 1 PoC trials is positive will the drug then be funded for stage 2 trials constituting the remaining $n - n_1$ PoC trials. The results of the stage 1 may be used to reallocate unspent funds from drugs that are unsuccessful in stage 1 to drugs that are successful in stage 1. Thus the latter drugs may be permitted more PoC trials than previously, based on the notion that their priority is increased due to successful performance in stage 1. A further refinement is that drugs may be prioritized for resources in stage 2 based on a Bayesian analysis with a prior estimate of p based on preclinical and Phase 1 data and a posterior estimate of p based on stage 1 results. We were interested in studying the properties of these arrays because they mirror practices in diversified pharmaceutical companies which are usually undertaken based on qualitative reasoning without quantitative optimization. We present some preliminary results here.

We modeled this problem as an urn problem, where each urn is a drug filled with green and red marbles, wherein green marbles represent indications for which the drug is truly active. We modeled a single urn (figure 3), where N = 15 and n = 5. The aggressive strategy entails $n_1 = n = 5$, ie all PoC trials in parallel as discussed above. The most conservative strategy entails $n_1 = 1$, and intermediate strategies $n_1 = 2$-4 were possible.

In the simulation with multiple urns (drugs), we assumed two urns with 15 marbles (PoC hypotheses) each, for a total of 30 marbles (figure 4). A total of 10 draws (PoC trials) are allowed. We draw n_1 marbles from each urn in stage 1. The allocation of the remaining $10 - 2n_1$ marbles depends on the results of stage 1. If only one urn yields at least one positive draw, that urn will receive all the remaining draws in stage 2. Under no circumstances will an urn that yields no positive PoC outcome in stage 1 receive further resources. If both urns yield at least one positive draw, the urn with the higher posterior estimate of p will be prioritized for the remaining draws. However, in this circumstance the inferior urn may receive leftover resources based on the total available resources and the number of marbles, the number of urns, and total number of allowed draws.

Figure 3:

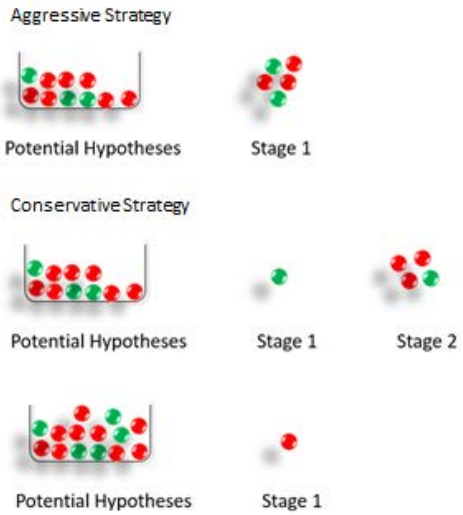

Figure 4:

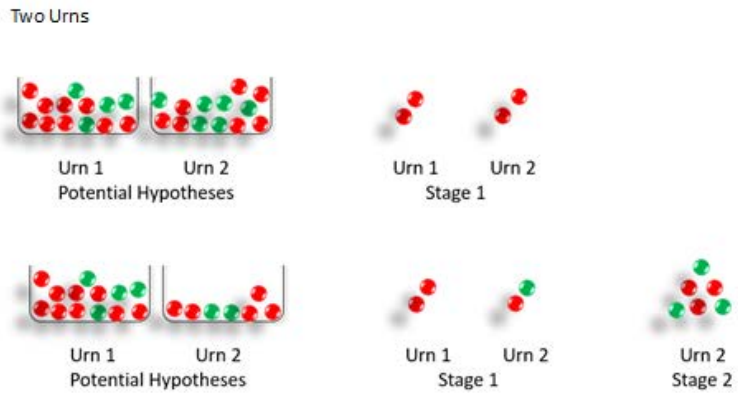

Clearly, the 2 stage approach in principle allows more intelligent resource distribution between drugs based on information from stage 1. However, it also delays development and may reduce throughput per unit time. In order to account for the effect of time, we implement two adjustments: 1. Discounting of benefits and costs occurring further in the future by an interest rate I. 2. A fixed overhead cost per unit time C_0 is applied, accounting for such non-development costs such as personnel, the research building and its electricity, etc. Thus, a conservative strategy will have higher fixed costs before benefits are seen.

The efficiency function is a generalization of the efficiency function given above. The benefits and costs are divided into two stages. The benefits and costs of the second stage

are discounted by two factors: 1. The probability that stage 2 will occur, and 2. The interest rate factor due to the delay relative to stage 1.

Detailed methods will be published elsewhere, but equations are given in the Appendix representing this problem. The likely distribution of green marble, N_G, in an urn is given by binom (N,p), while the number of green marbles in stage 1 is given by the hypergeometric distribution of (N, N_G, n_1). Further methods are presented briefly in the Appendix quantifying the observed number of green marbles in stage 1, which is not the same as the number of truly green marbles since false psoitives and false negatives will occur. Bayesian hierarchical modeling is performed to update the estimate of p based on stage 1 performance and the results of Stage 2 are simulated using this updated probability.

Figure 5 shows the efficiency, type III error, cost and benefit as a function of n_1 for C_0 = 0, delay of two years for stage 1, interest rate = 5%, and p = 0.5 for both urns. This figure gives a feel for the interplay of the parameters in this complex model. As expected, the conservative strategy is the most efficient in the absence of fixed costs because of the ability to update the initial estimate of p in allocating resources for stage 2. However, the efficiency improvement is only 1.5% and the absolute benefit is reduced by 40%, presumably due to the high probability of rejecting an effective drug based on only one PoC trial. Although the efficiency benefits of the conservative strategy are somewhat larger if the two urns have unequal values of p (not shown), the effect is still small compared to the dramatic reduction in throughput. This suggests that the optimum will change when fixed non-development costs are taken into account.

Figure 5:

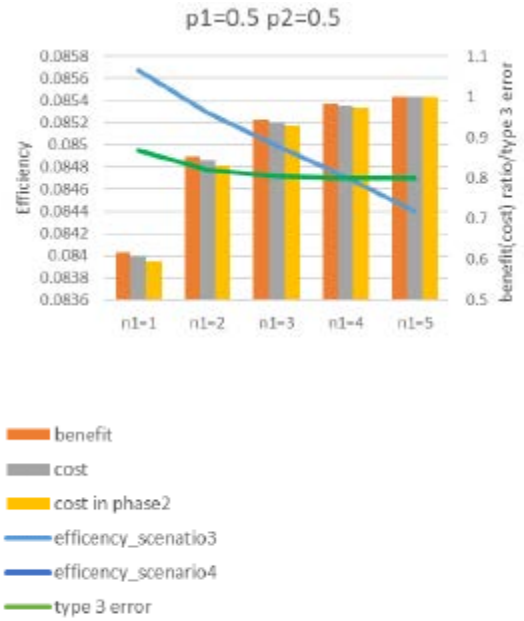

In the heat map on figure 6 (left side), we indeed see that with increasing fixed non-development costs C_0, the optimal efficiency is seen first at intermediate values of n_1, and

finally the aggressive strategy. Moreover, the conservative or aggressive strategy are always within 5% of the optimum, warranting a simplified optimality map shown on the right, given that in actual practice it may not be possible to determine the exact location of a portfolio on the heatmap.

Figure 6:

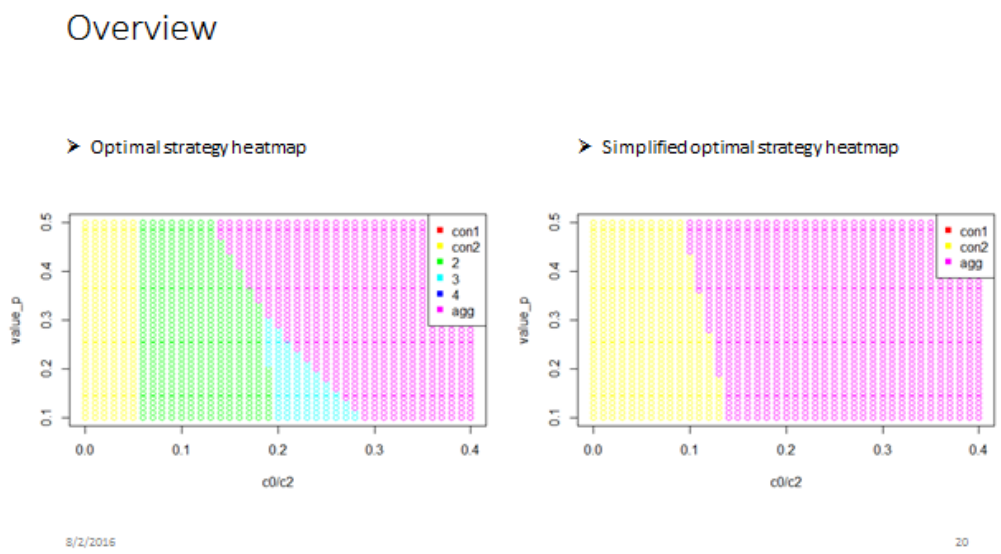

The optimal strategy represents a tradeoff between adaptively using information and maximizing throughput. As can be seen from Figure 6, the optimal strategy is relatively robust to the value of p, which is again fortunate, as this is difficult to estimate, and is largely determined by the value of the fixed non-development costs C_0. Only at low fixed costs is the conservative strategy optimal. Intermediate strategies may be optimal but are never dramatically superior to the extreme strategies under these conditions, thus potentially warranting a simplified optimality diagram.

The above analysis does not take into account other disadvantages of the conservative strategy such as diminished chance of first in class status. However, the above analysis also assumes that the prior and posterior estimates of p are reasonable and unbiased estimates of the outcome. The value of the conservative strategy may be much greater when this condition does not hold. Ongoing research is evaluating the various strategies when the true value of p and p' are very different from the estimates, for example when the drug developers believe drug A is better than drug B, when in fact drug B is better than drug A.

4. Future Directions

The work described herein considers efficiency from the point of view of the drug developer. We are also conducting ongoing research in which separate efficiency functions are developed from the viewpoint of different stakeholders. This research may illuminate scenarios where, for example, the incentives of sponsors and the public health need to be more carefully aligned [12].

Recently, we have been involved in a qualitative analysis of arrays of platform trials organized into clinical development pipelines [13]. The work described herein on arrays of PoC trials could be usefully applied to arrays of platform trials.

Finally, in the era of precision medicine portfolios should be considered to include not only putative therapies, but also putative predictive biomarkers as assets [7]. Methods to optimize such portfolios remain to be developed.

5. Acknowledgments

Cong Chen has been a strategic partner in the analysis, development, and application of the BCR for a decade.

We are indebted to Cong Chen and to Qiqi Deng for important comments on the analysis of 2-stage arrays of PoC trials

6. References

1. Chen, C. and Beckman, R. A., 2007. Optimal cost-effective designs of proof of concept trials and associated go-no go decisions. In *JSM Proceedings, Biometircs Section,* 394-399.
2. Chen, C., and Beckman, R.A. 2009. Optimal cost-effective Go-No Go decisions in late-stage oncology drug development. *Statistics in Biopharmaceutical Research,* 1: 159-169.
3. Chen, C., and Beckman, R.A., 2009. Optimal cost-effective designs of Proof of Conept trials and associated Go-No Go decisions. *Journal of Biopharmaceutical Statistics,* 19: 424-426.
4. Beckman, R.A., Clark, J., and Chen, C., 2011. Integrating predictive biomarkers and classifiers into oncology clinical development programs. *Nature Reviews Drug Discovery,* 10: 735-749.
5. Chen, C., and Beckman, R.A, 2014. Maximizing return on socioeconomic investment in Phase II proof-of-concept trials. *Clinical Cancer Research*, 20: 1730-1734.
6. Chen, C., Beckman, R.A. and Sun, L., 2015. Maximizing return on investment in Phase II proof-of concept trials. In *Optimization of Pharmaceutical R & D Portfolios,* edited by Zoran Antonijevic, Springer.
7. Beckman, R.A., and Chen C., 2015. Portfolio optimization of therapies and their predictive biomarkers. In *Optimization of Pharmaceutical R & D Portfolios,* edited by Zoran Antonijevic, Springer.
8. Song Y, and Chen, C, 2012. Optimal strategies for developing a late-stage clinical program with a possible subset effect. . *Statistics in Biopharmaceutical Research,* 4: 240-251.
9. Antonijevic, Z, 2016. The impact of adaptive design on portfolio optimization. *Therapeutic Innovation and regulatory Science,* 50: 615-619.
10. Lindborg, S.R., Persinger, C.C., Sashegyi, A., Mallinkrodt, C., and Ruberg, S., 2014. Statistical refocusing in the design of phase II trials offers promise of increased R & D productivity. *Nature Reviews Drug Discovery,* 13: 638-640.
11. Beckman, R.A., Schemmann, G. S. and Yeang, C.H. , 2012. Impact of genetic dynamics and single-cell heterogeneity on development of nonstandard

personalized medicine strategies for cancer. *Proceedings of the National Academy of Sciences USA,* 109: 14586-14591.

12. Ondra, T., Jobjörnsson, S., Beckman, R. A., Burman, C.F., König, F., Stallard, N., and Posch, M, 2016. Optimizing trial designs for targeted therapies. PLoS One 11:e0163726 .
13. Trusheim, M., Shrier, A. A.., Antonijevic, Z., Beckman, R. A., Campbell, R., Chen, C., Flaherty, K., Loewy, J., Lacombe, D., Madhavan, S., Selker, H., and Esserman. L.., 2016. PIPELINEs: Creating comparable clinical knowledge efficiently by linking trial platforms. Clinical Pharmacology and Therapeutics, in press..

7. Appendix 1

Underlying Model Assumption

- $N_G \sim$ Binomial (N, p)

---- N: Total number of PoC hypotheses for one drug; p: Estimated PoC trials positive rate; N_G: Number of positive PoC hypotheses

- $n_{1G} | N_G \sim$ Hypergeometric (N, N_G, n_1)

---- n1: Number of PoC trials in stage 1; n1G: Number of positive PoC trials in stage 1 (total number of trials permitted is n per um; n < N)

- $n_{1Gobs} | n_{1G} = n_{1Gobs\_TruePositive}$ (~ Binomial (n_{1G}, 1-β)) + $n_{1Gobs\_FalsePositive}$ (~ Binomial (n_1-n_{1G}, α)

---- n1Gobs: observed number of positive PoC trials in stage 1; α: Type I error; β: Type II error

- Bayesian hierarchical modeling: $P(N_G, n_{1G} | n_{1Gobs}) = P(n_{1Gobs} | n_{1G}) * P(n_{1G} | N_G) * P(N_G) / P(n_{1Gobs})$
- Posterior-p estimating: $E(p` = \frac{N_G}{N} / | n_{1Gobs}) = \frac{1}{N} \sum_{N_G} N_G * P(N_G, \sum n_{1G} | n_{1Gobs})$
- Efficiency Function (n1, p, C_0, I, y):

$$\frac{r^2 B_{2,1} + r^4 B_{2,2}}{\sum_y r^y C_0 + \sum_{y=0,1} r^y C_{2,1} / 2 + \sum_{y=2,3} r^y C_{2,2} / 2 + \sum_{y=2,3,4} r^y C_{3,1} / 3 + \sum_{y=4,5,6} r^y C_{3,1} / 3} \quad ; \quad (where\ r = \frac{1}{1+I})$$

---- I: Interest rate; y: Number of years delay in executing Stage 1; C_0: Fixed non-development cost per year;

B_{ab}: Estimated benefit from phase a trials in stage b; C_{ab}: Estimated cost for phase a trials in stage b or originating from stage b

8/2/2016 18