PHARMALEX

Bayesian Adaptive Phase I-II Designs for Evaluating Safety and Efficacy in Dual-Agent Oncological Drug Development

Bradley P. Carlin, PhD

with special thanks to my PharmaLex coauthors

Lira Pi, PhD, James Willard, MS, Bruno Boulanger, PhD, Elisabeth Rouits, PhD, and Maud Hennion, MS

Prepared for BAYES2023 meeting, October 25th, 2023

Disclaimer

- The information provided during this presentation does not constitute legal advice. PharmaLex, and its parent Cencora, strongly encourage the audience to review available information related to the topics discussed during the presentation and to rely on their own experience and expertise in making decisions related thereto. Further, the contents of this presentation are owned by PharmaLex and reproduction of the slides used in today's presentation is not permitted without consent of PharmaLex.

- *Conflict of interest statement:* Bradley P. Carlin is a member of the scientific organizing committee for BAYES2023.

Introduction

- Traditional approach to Phase I testing in oncology:
  - Phase 1: Safety only (3+3, CRM, BOIN, etc.) to determine a recommended Phase 2 dose (RP2D), typically the maximum tolerated dose (MTD)
  - Phase 2: Preliminary Efficacy with safety monitoring
- Approach works well with traditional cytotoxic agents, but not with more modern therapies (e.g. vaccines)
  - Efficacy may rise quickly with dose, and plateau long before MTD is reached
- New approach proposed in **FDA Project Optimus**:
  - "Multiple dosages should be compared in a clinical trial(s) designed to assess activity, safety, and tolerability… in a dose-finding trial"
  - "A recommended trial design to compare these dosages is a randomized, parallel dose-response trial."
    - "An adaptive design to stop enrollment of patients to one or more dosage arms of a clinical trial following an interim assessment of efficacy and/or safety could be considered."
  - "The analysis plan should specify a multiple-testing procedure which accounts for testing multiple treatments versus a control as well as any interim assessments after which an inferior arm is dropped."

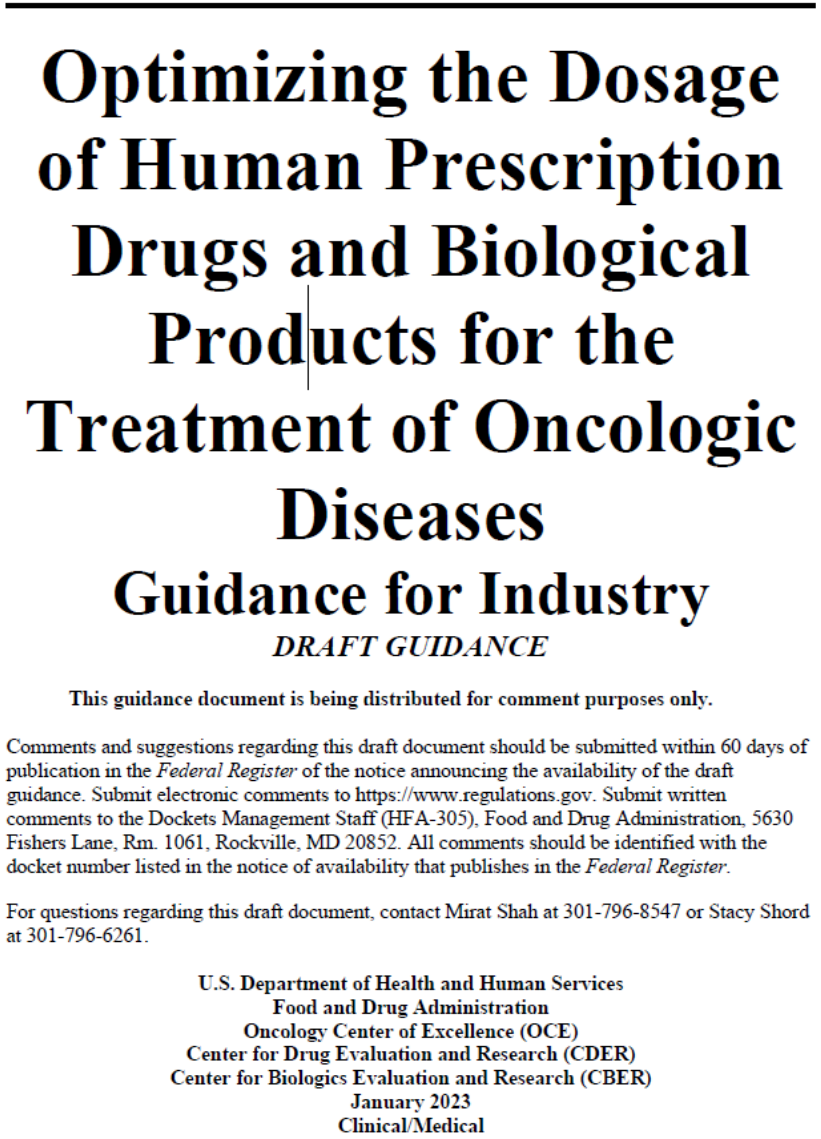

Optimizing the Dosage of Human Prescription Drugs and Biological Products for the Treatment of Oncologic Diseases

Guidance for Industry

DRAFT GUIDANCE

This guidance document is being distributed for comment purposes only.

Comments and suggestions regarding this draft document should be submitted within 60 days of publication in the *Federal Register* of the notice announcing the availability of the draft guidance. Submit electronic comments to https://www.regulations.gov. Submit written comments to the Dockets Management Staff (HFA-305), Food and Drug Administration, 5630 Fishers Lane, Rm. 1061, Rockville, MD 20852. All comments should be identified with the docket number listed in the notice of availability that publishes in the *Federal Register*.

For questions regarding this draft document, contact Mirat Shah at 301-796-8547 or Stacy Shord at 301-796-6261.

U.S. Department of Health and Human Services
Food and Drug Administration
Oncology Center of Excellence (OCE)
Center for Drug Evaluation and Research (CDER)
Center for Biologics Evaluation and Research (CBER)
January 2023
Clinical/Medical

Takeaways from Project Optimus guidance

- Phase I needs to consider both safety and efficacy, and to search not for an RP2D, but for a recommended dose *range* (RDR) to advance to Phase 2
- Phase 2 needs to be a randomized comparison of a few doses from the RDR:
  - the maximum tolerated dose (MTD), if reached
  - the lowest effective dose (LED), i.e. the safest dose that achieves a desired efficacy level
- Oncology drugs are often given in two-drug **combinations,** meaning we must search a two-dimensional space, and the RDR becomes a recommended dose ***region***
  - Consider a new drug (Drug 1), which can be given alone or in combination with an existing drug (Drug 2) whose optimal dose as a monotherapy is already approved for use. We may now compare placebo/SoC to:
    - The best 2 doses of Drug 1 alone
    - These same doses when combined with the existing approved Drug 2 dose
    - The best 2 combination therapies where both doses are unrestricted
- Suggests an algorithm for Phase 1 dose-finding:
  - First find the monotherapy RDR (lies on the d_1 axis here)
  - Then find the combination therapy RDR by searching the entire space
    - Here, the optimum lies at $(d_1, d_2) = (1, 0.5)$
- Best way to search this space? (Doses no longer ordered…)
- What statistical model shall we use? (Worth modeling correlation?...)

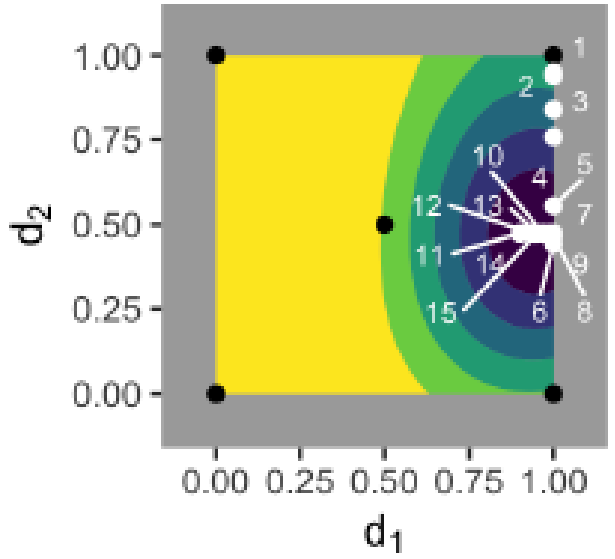

Model and Algorithm for Bivariate Dose-Finding

- Efficacy and safety can be modeled jointly (bivariate normal, copula, etc); c.f. Mu et al. (2021)
  - Indeed this seems most sensible (more efficacious drugs will sometimes be less safe)
  - BUT this typically requires MCMC methods, and results are often similar to those from models that simply consider safety and efficacy separately (Guo and Yuan, 2023; many others…)
- In the case where both the safety and efficacy outcomes are continuous, Willard et al. (2023) use a Gaussian Process (GP) approximation to the bivariate response surface
  - Could be efficacy alone, or a utility function that trades off efficacy and safety
- The paper then uses Bayesian Optimization (BayesOpt; Garnett, 2023), a derivative-free maximizer that selects the next design point to evaluation using an acquisition function
  - trades off exploitation (regions we think have good values)
  - and exploration (regions we haven't visited much yet)
- Approach is also extended to personalized dose-finding
  - Here, a single binary covariate Z determines 2 response surfaces
  - Easily extended to $P > 1$ discrete covariates
- Operating characteristics can be checked by simulation
  - Euclidean distance between true and estimated d_{opt}
  - Expected root MSE
- *Want full details?* Attend James Willard's talk, 10 am Friday!

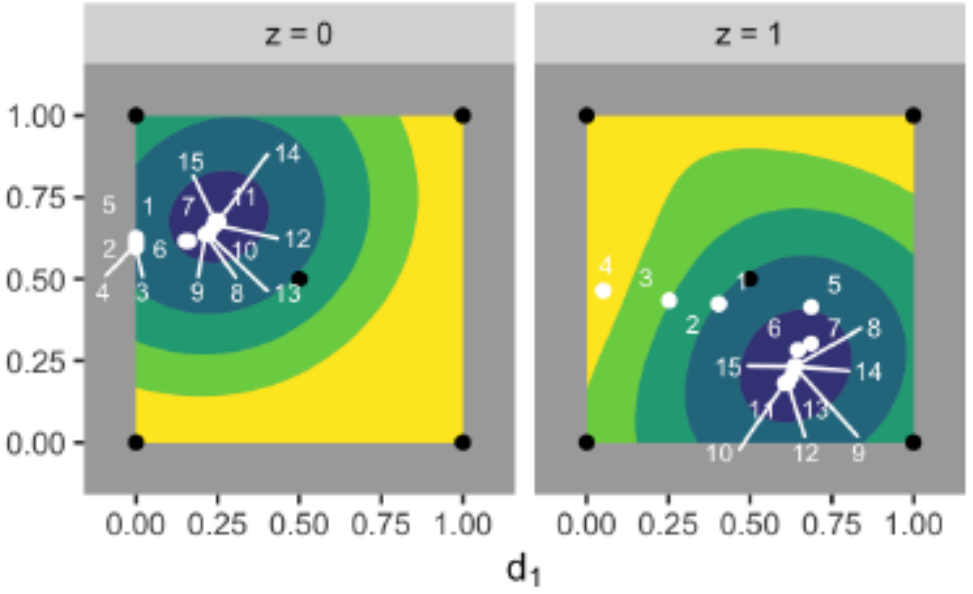

Our Application: Combination therapy for Prostate Cancer

- *Our setting:* the safety outcome is binary (DLT/no DLT), but the efficacy outcome is continuous (absolute improvement in PSA score)
- Client also wishes to compare two dosing regimens:
  - **Regimen 1 ("4/3")** assumes the drug is given 4 days per week, followed by a 3-day break
    - Doses are pre-specified as 100 mg, 150 mg, 200 mg, and 250 mg per day, hence total doses are (400, 600, 800, 1000) mg per week.
    - Our doses are standardized by the max dose (1000 mg/week) to become **(0.40, 0.60, 0.80, 1)**
  - **Regimen 2 ("5/2")** instead assumes 5 days on, 2 days off per week
    - Now the weekly dose-levels become (500, 750, 1000) mg per week, so here the standardized doses are **(0.50, 0.75, 1)**
  - Assume cohorts of size 3, and do not permit more than 3 cohorts at any one dose
- Design assumptions given two parallel regimens:

Feature	Regimen 1 (4/3)	Regimen 2 (5/2)
# of simulation data sets	100	100
Standardized dose-level	(0.40, 0.60, 0.80, 1.00)	(0.50, 0.75, 1.00)
# of patients per cohort	3	3
Max. # of cohorts	8	6

Modeling Example: Safety

- Statistical Model:

Let $Y_{ij} = 1$ if patient i taking dose j experienced a DLT, and 0 otherwise. Then

$$Y_{ij} \sim Bernoulli(\theta_j)$$

Assume a linear logistic response model

$$logit(\theta_j) = \alpha_0 + \alpha_1 d_j \,,$$

Once we specify the target (highest acceptable) toxicity level (TTL), this induces a posterior on d_U,

$$d_U = [logit(TTL) - \alpha_0]/\alpha_1$$

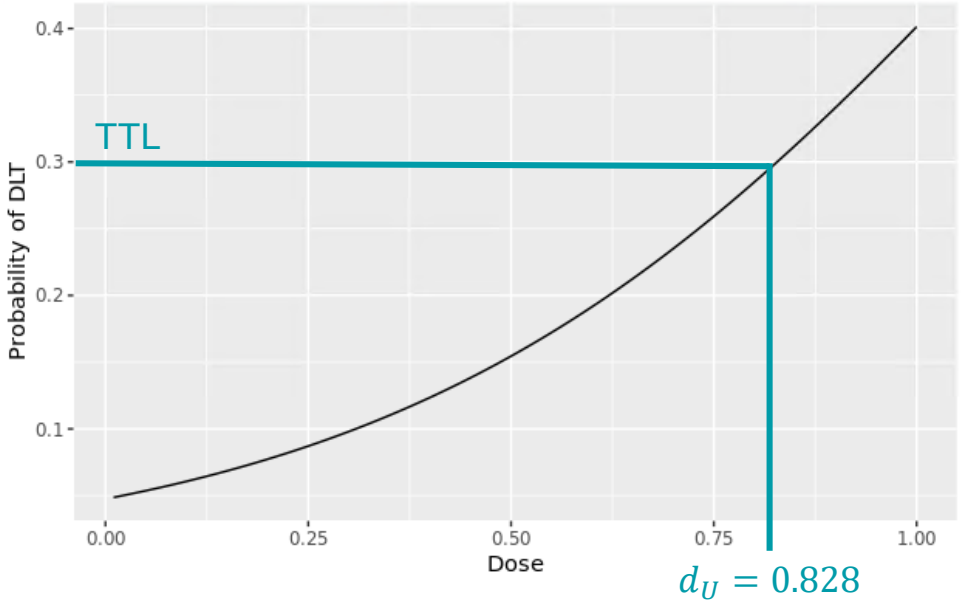

Parameter	Assumption
α_0	-3
α_1	2.6
TTL	0.3

induces

Dose	Pr(DLT | regimen 1)	Pr(DLT | regimen 2)
$d_1 = 0.40$	0.1235	
$d_2 = 0.50$		0.1545
$d_3 = 0.60$	0.1915	
$d_4 = 0.75$		0.2592
$d_5 = 0.80$	0.2850	
$d_U = 0.828$	0.30 (= TTL)	
$d_6 = 1.00$	0.4013	

Modeling Example: Efficacy

- Statistical model
  - Let Z_{ij} be the absolute improvement (drop) in PSA score between baseline (BL) and 6 months for patient i taking dose j.
  - Assume

$$Z_{ij} \sim N(\mu(d_j), \sigma^2), \qquad where\ \mu(d_j) = \beta_0 + \beta_1 d_j + \beta_2 d_j^2$$

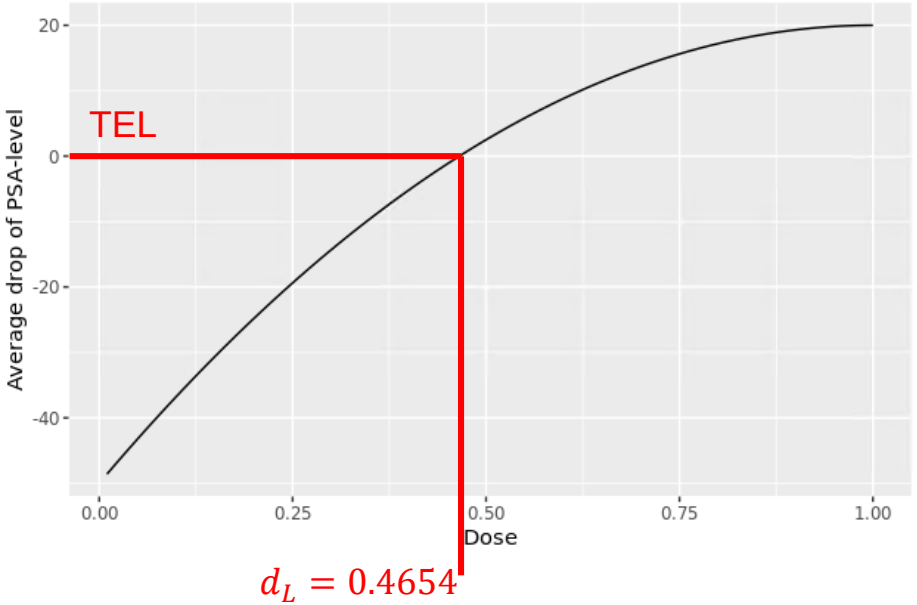

By using the roots of quadratic equation given the target efficacy level (TEL),

$$d_L = \frac{-b + \sqrt{b^2 - 4ac}}{2a}$$

Where $a = \beta_2;\ b = \beta_1; c = \beta_0\ - TEL.$

Assumptions:

Parameter	Assumption
β_0	-50
β_1	140
β_2	-70
σ	10
TEL	0

induces

Dose	Avg. PSA drop in regimen 1	Avg. PSA drop in regimen 2
$d_1 = 0.40$	-5.2	
$d_L = 0.4654$	0.0 (= TEL)	
$d_2 = 0.50$		2.5
$d_3 = 0.60$	8.8	
$d_4 = 0.75$		15.6
$d_5 = 0.80$	17.2	
$d_6 = 1.00$	20.0	

Defining the RDR (Recommended Dose Range)

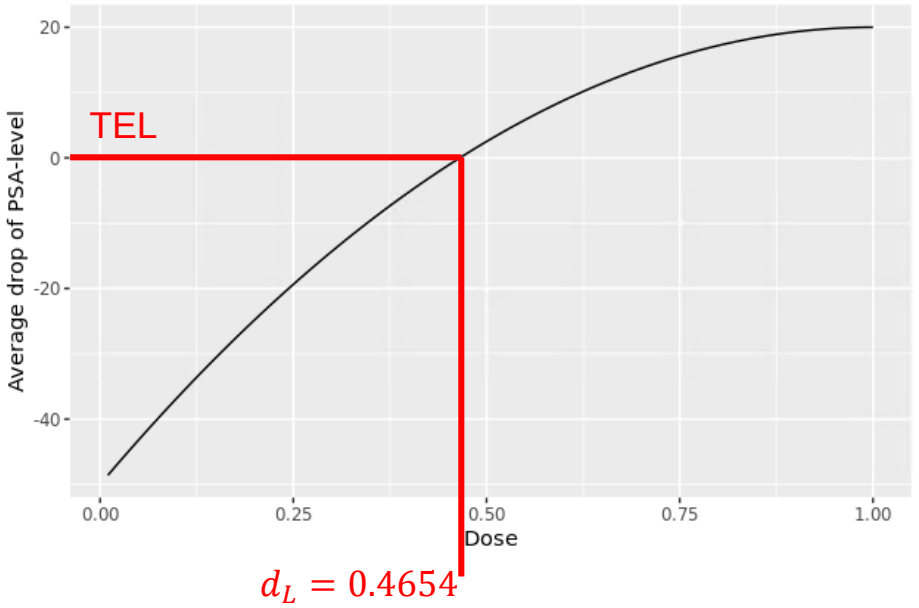

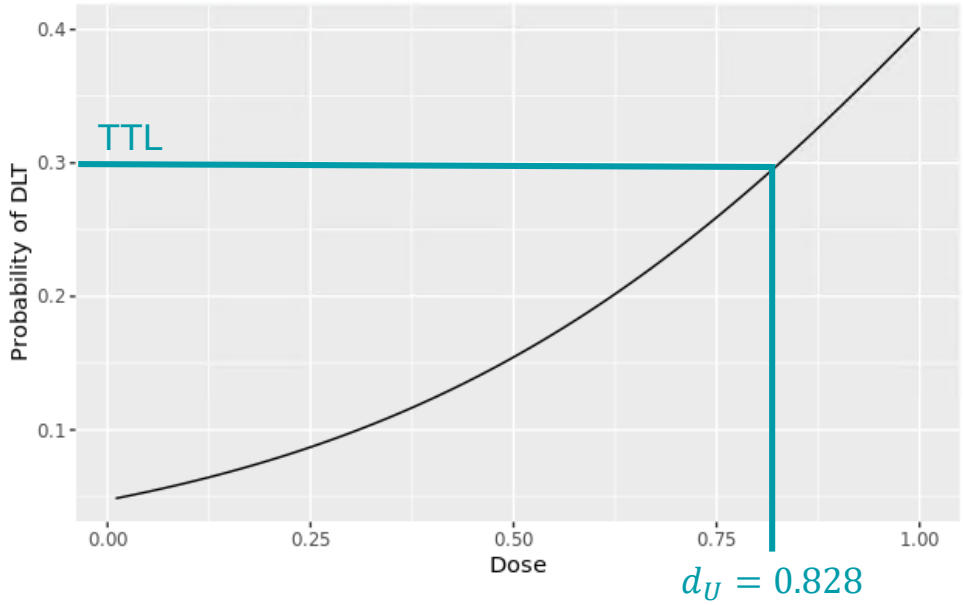

- By our assumptions above, the hypothetical RDR should consist of the doses within this range, ie.,

$$RDR = \{0.5, 0.6, 0.75, 0.8\} = \{d_2, d_3, d_4, d_5\}$$

Bivariate CRM Approach to Monotherapy Dose-Finding

- General Approach:
  - Start at the lowest dose, d_1
  - Find d_U, the highest safe dose (i.e., that has toxicity at most equal to the TTL, 0.3)
  - "Backfill" patients at lower doses in order to find d_L, the lowest effective dose (i.e., that has efficacy at least equal to the TEL, 0)
  - Do NOT permit "dose skipping" (i.e., cannot jump from d_1 to d_3)
  - Assume maximal total sample sizes of 24 for Regimen 1 and 18 for Regimen 2 (i.e., average of 6 patients per dose)
  - Do not assign more than 9 patients to any one dose
- Three stopping rules are necessary:
  - Stop the search for d_U when the width of its 95% BCI is less than 0.3 $(= (d_6 - d_1)/2)$
  - Stop the search for d_L when the width of its 95% BCI is less than 0.1 (efficacy endpoint is "easier" to learn about since it is continuous, not binary)
  - Stop both searches when we hit the maximum sample size (24 or 18), even if the interval width stopping rules has not engaged
- To judge our design, we need to simulate its operating characteristics
  - Probability of correct identification of the true RDR
  - Average trial length / average sample size
- True scenarios to be investigated:
  - RDR = $(d_1, d_2), (d_1, d_4), (d_1, d_6), (d_2, d_2), (d_2, d_5)$, etc.

Simulated Distributions of $\hat{d}_L$ and $\hat{d}_U$

- Quartiles for regimen 1 (4/3) are

	Min.	25th	Median	75th	Max.
$\hat{d}_U$	0.7912	0.8194	0.8394	0.8542	0.8890
$\hat{d}_L$	0.4087	0.4460	0.4588	0.4752	0.5032

- Quartiles for regimen 2 (5/2) are

	Min.	25th	Median	75th	Max.
$\hat{d}_U$	0.7840	0.8188	0.8380	0.8495	0.8865
$\hat{d}_L$	0.4104	0.4464	0.4566	0.4762	0.5177

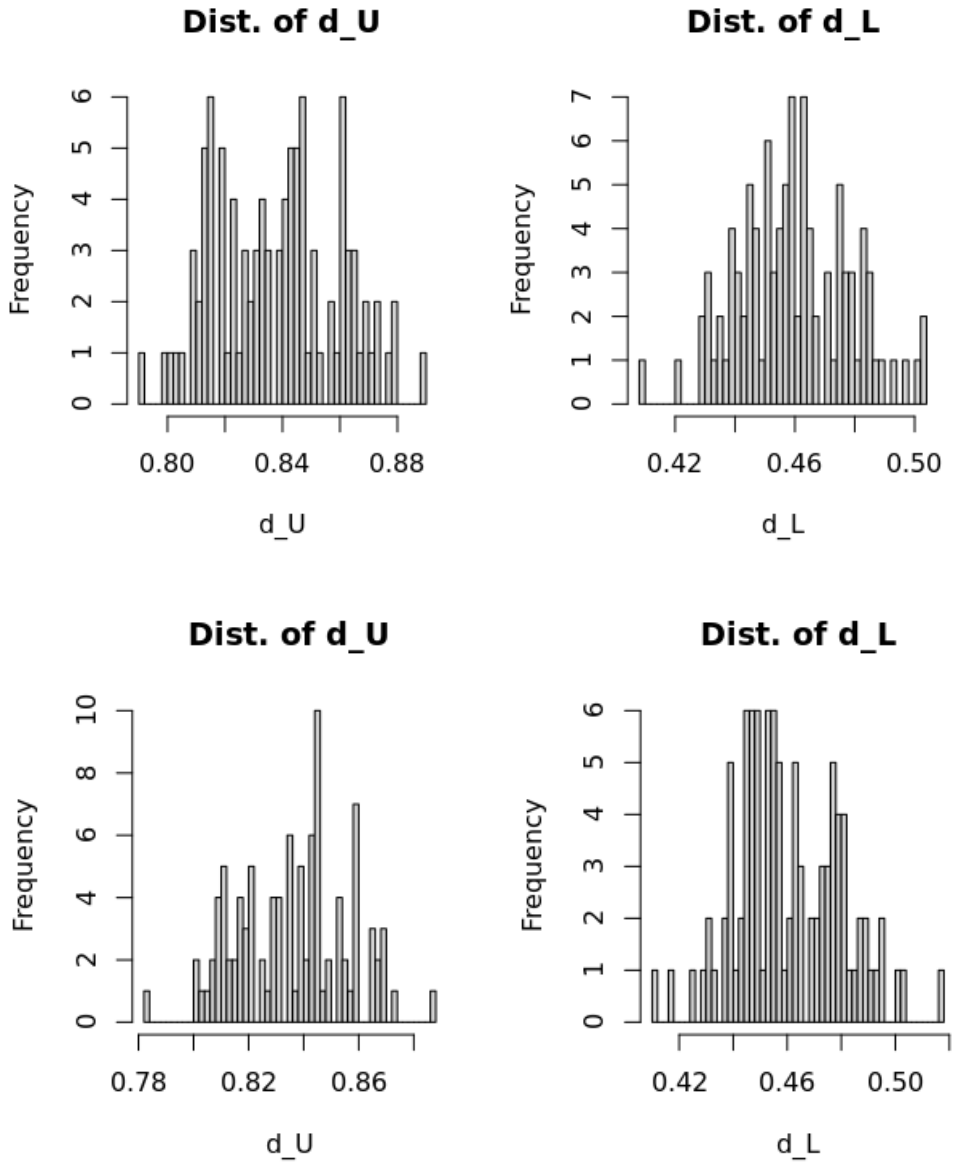

- Given the simulation assumptions, true $(d_L, d_U) = (0.4654, 0.8280)$ are quite similar to the medians of $(\hat{d}_L, \hat{d}_U) = (0.4588, 0.8394)$ for regimen 1 and the other medians of $(\hat{d}_L, \hat{d}_U)$ = (0.4566, 0.8380) for regimen 2.

- In this example,
  - $d_1 = 0.40$ is safe but not effective
  - $d_6 = 1.00$ is effective but not safe

- Thus on average, the RDR does include the correct four doses, $\{d_2 = 0.5, d_3 = 0.6, d_4 = 0.75, d_5 = 0.8\}$.

Probability of Correct Identification of the RDR ("Power Analysis")

- Contingency table of dose selection:
  - d_U_false includes two error cases: d_U < 0.8 (doses of 0.8 or lower are wrongly dropped) or d_U > 1.0 (dose $d_6 = 1.0$ is incorrectly included)
  - So d_U_true when $d_U \in [0.8, 1.0)$
  - d_L_false includes two error cases: d_L > 0.5 (doses of 0.5 or more are wrongly dropped) or d_L < 0.4 (dose $d_1 = 0.4$ is incorrectly included).
  - So d_L_true when $d_U \in (0.4, 0.5]$
- In this simulation, our algorithm did very well:
- For regimen 1 (4/3):

	d_U_true	d_U_false	Total
d_L_true	95	2	97
d_L_false	3	0	3
Total	98	2	100

```
> d_tab %>% filter(u_yes == 0 | l_yes == 0)
       d_u2 R2      d_l2 R2_dl u_yes l_yes
1 0.7911529  3 0.4424998     3     0     1
2 0.7996996  3 0.4457929     3     0     1
3 0.8628734  8 0.5013415     8     1     0
4 0.8697908  8 0.5029821     5     1     0
5 0.8775575  8 0.5031778     8     1     0
```

- For regimen 2 (5/2):

	d_U_true	d_U_false	Total
d_L_true	96	1	97
d_L_false	3	0	3
Total	99	1	100

```
> d_tab %>% filter(u_yes == 0 | l_yes == 0)
       d_u2 R2      d_l2 R2_dl u_yes l_yes
1 0.7839571  3 0.4439707     3     0     1
2 0.8543600  6 0.5176643     6     1     0
3 0.8683984  6 0.5029001     5     1     0
4 0.8594754  6 0.5011979     6     1     0
```

Average Trial Length (Number of Cohorts Required)

- Frequency table of number of cohorts needed to find the RDR among 100 fake data sets, Regimen 1 (4/3):

month	1	2	3	4	5	6	7	8
Safety	0	1	10	7	9	5	3	65
Efficacy	1	3	12	15	16	12	12	29

- Frequency table of number of cohorts needed to find the RDR among 100 fake data sets, Regimen 2 (5/2):

month	1	2	3	4	5	6
Safety	0	4	12	8	4	72
Efficacy	1	6	14	15	15	49

- Total Phase I trial duration in months (4 weeks ≈ 1 month; safety endpoint in 4 wks, efficacy endpoint in 12 wks)

month	4	5	6	7	8	9	10
Regimen 1	1	11	7	10	30	12	29
Regimen 2	4	12	20	15	49	n/a	n/a

Avg trial length (months)
8.09
6.93

Distribution of Course Length

month	4	5	6	7	8	9	10
Regimen 1	1	11	7	10	30	12	29
Regimen 2	4	12	20	15	49	n/a	n/a

Regimen 1 (4/3)

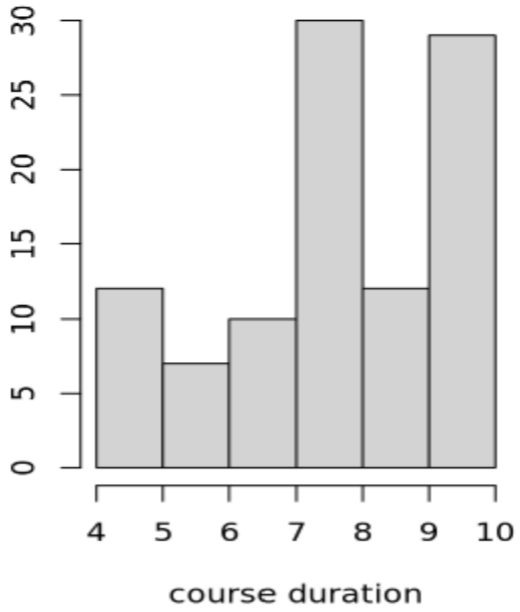

Regimen 2 (5/2)

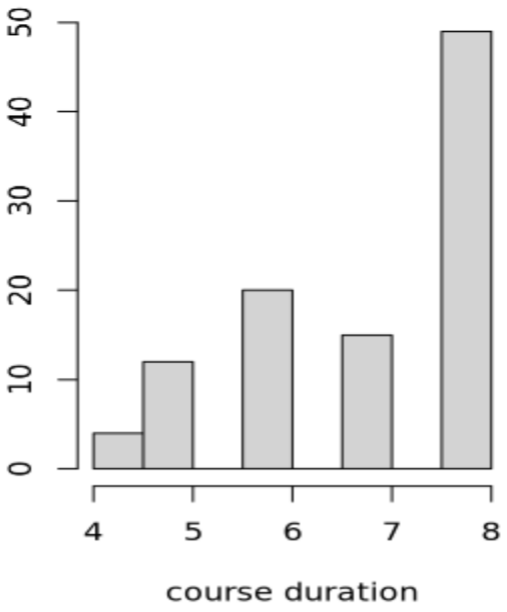

Other things we'd want to investigate:

- Different true arrangements of the RDR (e.g., $(d_1, d_6), (d_2, d_2)$, etc.)
- Settings where our initial guesses for α_1 and/or the βs are incorrect (e.g., centered on the wrong values)

Next Step and Conclusion

- Moving from monotherapy RDR (range) to combination therapy RDR (region)
  - *Could* search the bivariate dosing space after the monotherapy RDR was found, BUT
    - Time consuming (16 months on average instead of 8?)
    - Would need to be done for each candidate combination drug (it turns out there is not 1, but 3 or 4)
    - Would require a separate IND if we depart from the approved doses of the combo drugs
- Designing the accompanying randomized Phase II study
  - Trial would compare the two most promising monotherapy doses, 3-4 combo doses (combos at approved levels + lowest dose in the RDR)
  - Trial needs a Placebo/Standard of Care arm, but investigators do not agree on its definition
  - Primary endpoint: PSA improvement at 12 months
  - Secondary endpoints: OS, PFS (PSA- or radiographically-based), SAEs/other safety endpoints
  - Need to collect data that will inform Phase III design (e.g., pbo + combo vs d_2 + combo)
  - Helpful adaptive steps: Early look(s) for early stopping for futility / safety (esp for combo arms) / early efficacy?
  - FDA guidance makes clear multiple testing adjustments needed to account for the multiple doses tested
- Conclusion: Bayesian methods can help sort out conflicting goals and provide a sensible "roadmap" for satisfying new Project Optimus requirements for Phase 1-2!
  - Approach likely to be helpful/needed in *non-oncological* areas as well!

References

1. Carlin, B.P. and Louis, T.A. (2009). *Bayesian Methods for Data Analysis, 3rd ed.* Boca Raton, FL: CRC Press.
2. FDA (2023). *Optimizing the Dosage of Human Presciption Drugs and Biological Products for the Treatment of Oncological Diseases: Guidance for Industry.* Food and Drug Administration, Center for Drug Evaluation and Research (CDER) and Center for Biologics Evaluation and Research (CBER), Rockville, MD.
3. Garnett, R. (2023). *Bayesian Optimization.* Cambridge, UK: Cambridge University Press.
4. Guo, B. and Yuan, Y. (2023). DROID: dose-ranging approach to optimizing dose in oncology drug development. *Biometrics,* 1-13. https://doi.org/10.1111/biom.13840.
5. Mu, R., Xu, J. ,Tang, R., Kopetz, S., and Yuan, Y. (2021). A Bayesian phase I/II platform design for co-developing drug combination therapies for multiple indications. *Statistics in Medicine*, **41**, 374—389.
6. Price, D. and Scott, J. (2021). The US food and drug administration's complex innovative trial design pilot meeting program: progress to date. *Clinical Trials,* **18(6)**, 706-710. https://doi.org/10.1177/17407745211050580