Early Phase Studies in Industry

Early Phase Studies

- Early phase study design will depend on a multitude of factors (e.g.; disease, study population, drug, regulatory, etc…)
- For example, a new delivery method for an old drug may require a bioequivalence study (cmax and auc)
- A new drug (or at least a new drug for the study population) will go through some sort of dose finding/ranging

Early Phase Studies

- Many early phase studies employ Bayesian statistical methods
- A lot of folks have not been properly trained in Bayesian methods

The Teaching of Bayesian Methods

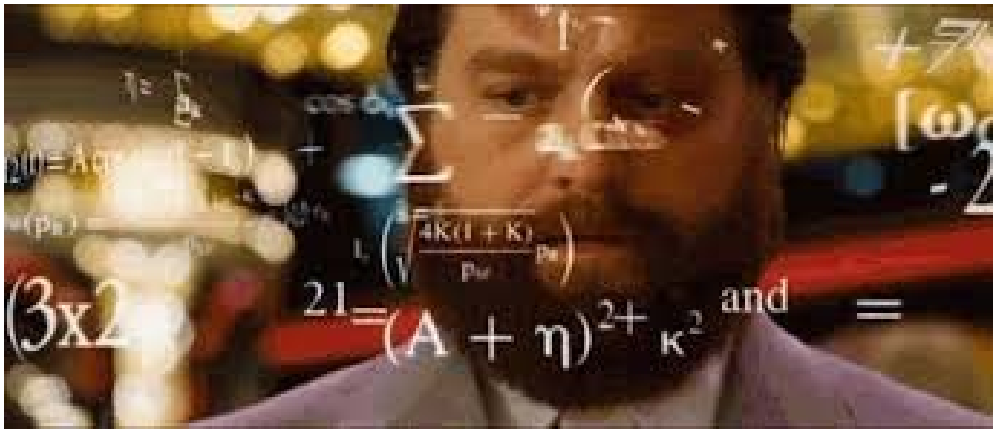

- There are probably introductory statistical courses that do not cover Bayesian methods (or at least there *were* intro courses that did not cover this topic)
- Yet Bayesian methods grow ever more common, especially in early phase studies
- There can be hesitancy to implement these methods because they appear ‘new’ or ‘complicated’

Overview

- For this talk, we will walk through:
  - A real-world example of Bayesian thinking
  - Some “simple” exercises that could be used in early phase studies:
    - An example of a frequentist confidence interval
    - An example of a Bayesian credible interval
    - An example of a Bayesian posterior probability against a threshold
  - A few additional applications used in early phase trials

Bayesian Inference in the Real World

- Bayesian inference is the reallocation of credibility across possibilities
- Suppose we were to step outside one morning to find that the sidewalk is wet…
- Why is the sidewalk wet?

Maybe…

- …it has recently rained?
- …someone watered the grounds nearby?
- …a new spring had erupted from underground?
- …a sewage pipe has broken?
- …a passerby spilled a drink?

Prior Credibility

- If all we know up to this point is that the sidewalk is wet, then all of the previous possibilities will have some level of credibility based on our previous knowledge
- However, continuing on our journey outside, we may collect additional observations…

If we observe that…

- …the sidewalk is wet for as far as we can see
- …the trees are wet as well
- …the cars are wet, too
- We may then re-allocate credibility to the hypothetical cause of rain

If we instead observe that…

- …the wetness is localized to a small area
- …there is an empty drink cup nearby
- We may instead re-allocate credibility to the hypothetical cause of a spilled drink

If we instead observe that…

- …Clark Griswold is nearby
- …there is a house with an obscene amount of lights nearby
- We may actually allocate credibility to the possibility of Cousin Eddie's dirty work
- This re-allocation of credibility from areas of prior belief to new areas based on accumulating data is the essence of Bayesian inference

Example

- Suppose that we collect treatment-emergent adverse event (TEAE) data
- Each of 30 subjects will either have at least one TEAE or they will have zero TEAEs
- Suppose that 10 out of the 30 subjects observe at least one TEAE

A 95% Confidence Interval

- We can use a formula (and a normal approximation):

$$\hat{p} \pm z_{1-\frac{\alpha}{2}} \sqrt{\frac{\left(\hat{p} * (1-\hat{p})\right)}{n}}$$

$$= 0.33 \pm 1.96 \sqrt{\frac{0.33 * (1-0.33)}{30}}$$

$$= (0.16, 0.50)$$

- Interpretation: If we repeated this experiment 100 times, we would expect 95% of the confidence intervals to contain the true value of the probability of a TEAE

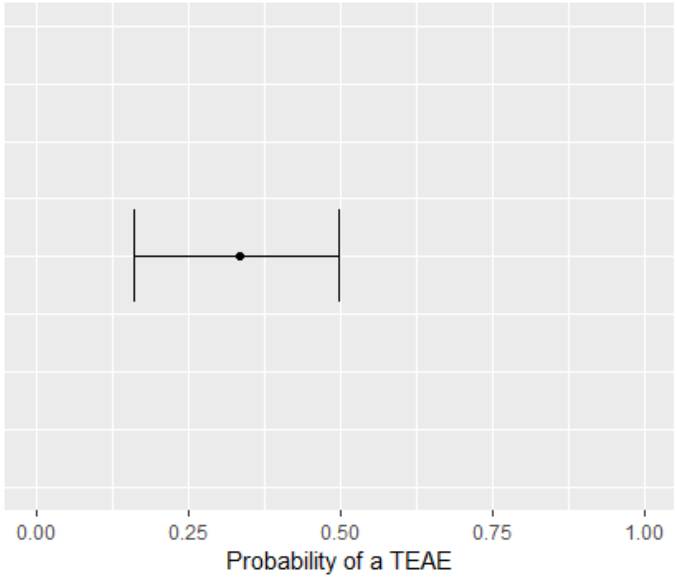

Bayesian Analysis and the Beta-Binomial

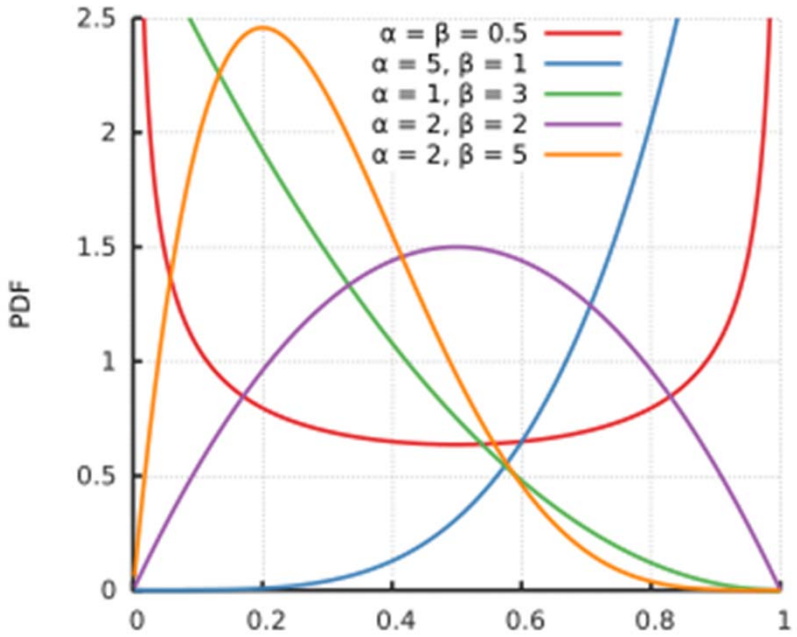

- These are beta distributions for various values of the distributions two parameters, alpha and beta
- Alpha (sometimes called shape1) can loosely be interpreted as “the number of events that occurred”
- Beta (sometimes called shape2) can be loosely interpreted as “the number of trials in which an event did not occur”
- Therefore, every time we observe an event, we add 1 to alpha and every time we have a subject that does not observe the event, we add 1 to beta
- To get any point along any of these lines, use “dbeta(x, shape1=alpha, shape2=beta)” in R
- For example: dbeta(0.5, shape1=2, shape2=2) will give the maximum point on the purple line to the left.

An Example of a Bayesian Approach

- We will need to pick our prior belief of the probability of a TEAE
- One potential choice (that may actually seem pretty counterintuitive, but can actually work out nicely) is to choose:
  - Alpha = 1
  - Beta = 1
- So our prior distribution is a BETA(1, 1)

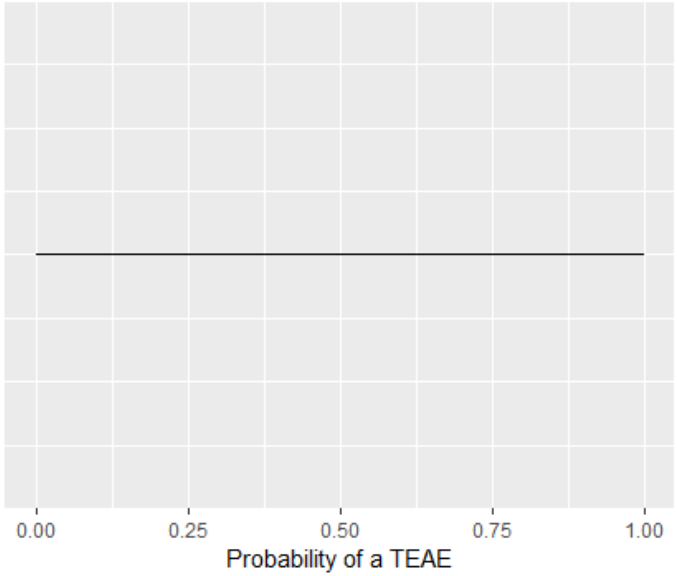

Suppose the first subject did not have an event…

- Alpha = 1 + 0
- Beta = 1 + 1
- So our posterior distribution is now a BETA(1, 2)

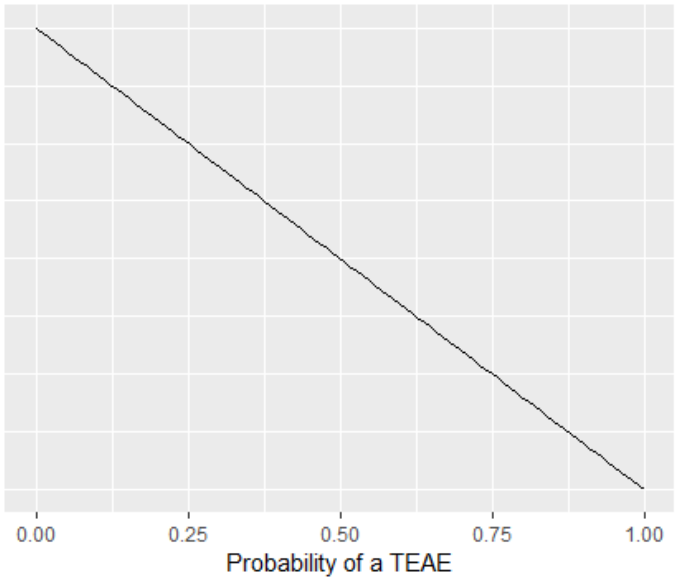

Suppose that the second subject did not have an event either…

- Alpha = 1 + 0 + 0
- Beta = 1 + 1 + 1
- Our posterior distribution is now BETA(1, 3)

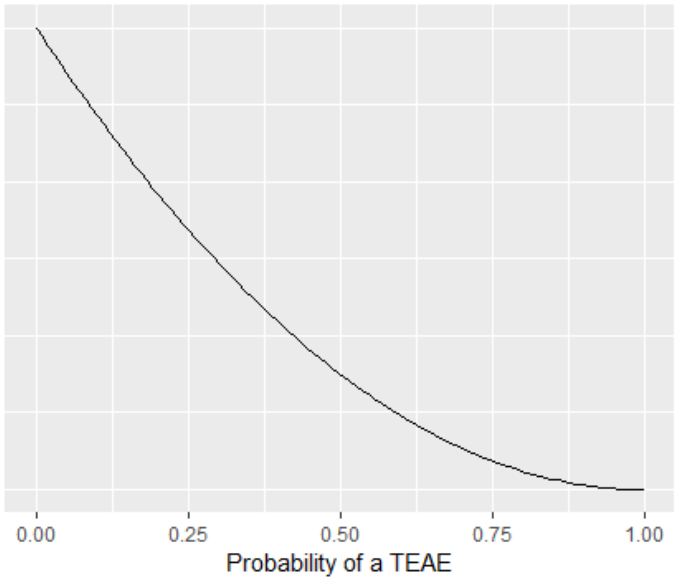

After all 30 subjects (with 10 events)

- Alpha = 1 + 10
- Beta = 1 + 20
- Our posterior distribution is now BETA(11, 21)

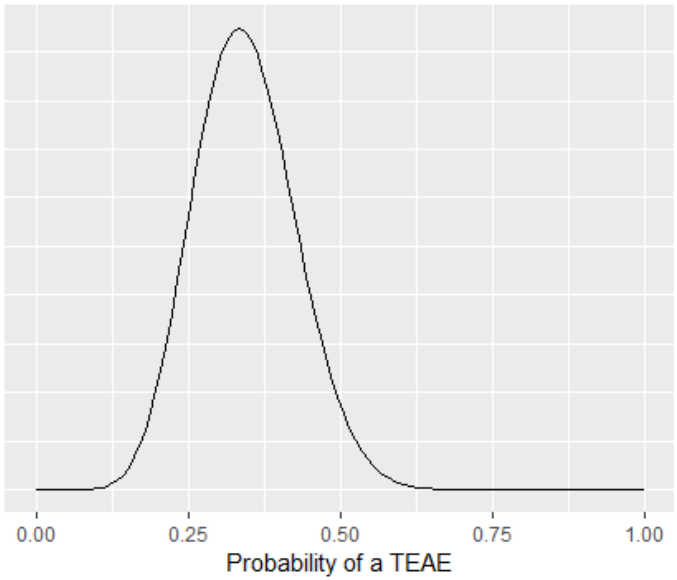

Building a 95% Credible Interval

- One way to obtain a 95% credible interval would be to take the 2.5% and 97.5% percentiles of the BETA(11, 21) distribution:

$$(0.19, 0.51)$$

- The above is done in R by the following two commands:
  - qbeta(0.025, 11, 21)
  - qbeta(0.975, 11, 21)
- This is pretty close to the 95% confidence interval:

$$(0.16, 0.5)$$

- The interpretation is: 95% of the posterior probability of a TEAE lies between 0.19 and 0.51.

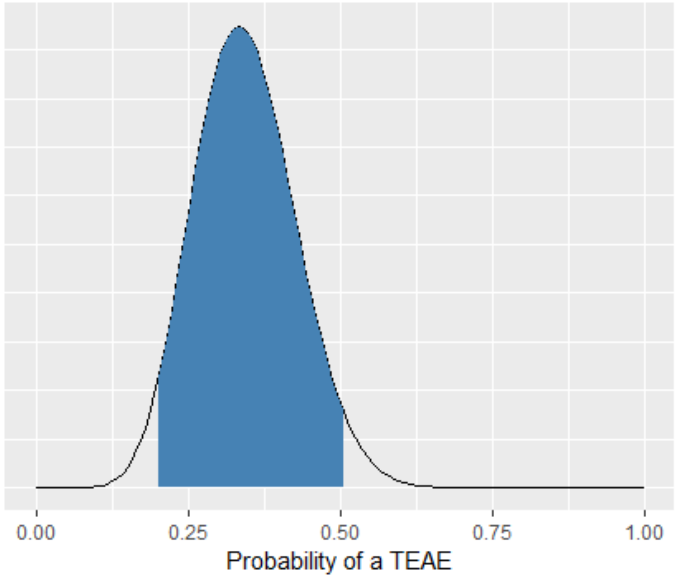

Thresholds

- If there were some special threshold, we could speak to the probability of being above or below the threshold
- For example, if 40% were a specific threshold for our disease of interest, both the frequentist confidence interval and the Bayesian credible interval contain 0.4
- However, it is quite natural, in the Bayesian framework, to say “75% of the posterior probability lies below 0.4”
- The above is accomplished with the R command:
  - pbeta(0.4, 11, 21)

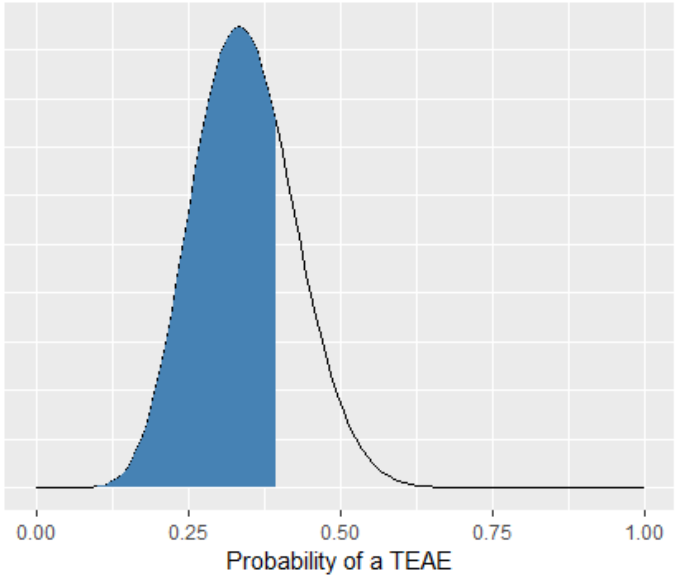

CRM Example: Choose Dose Levels and Prior Probabilities

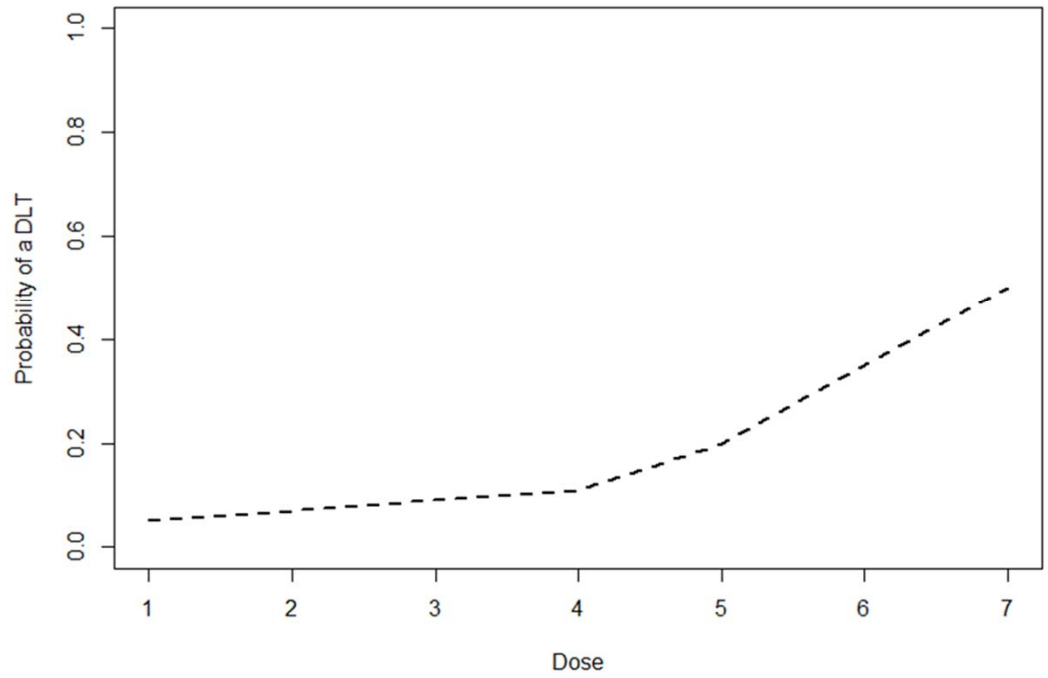

CRM Example: Choose Target Toxicity Level

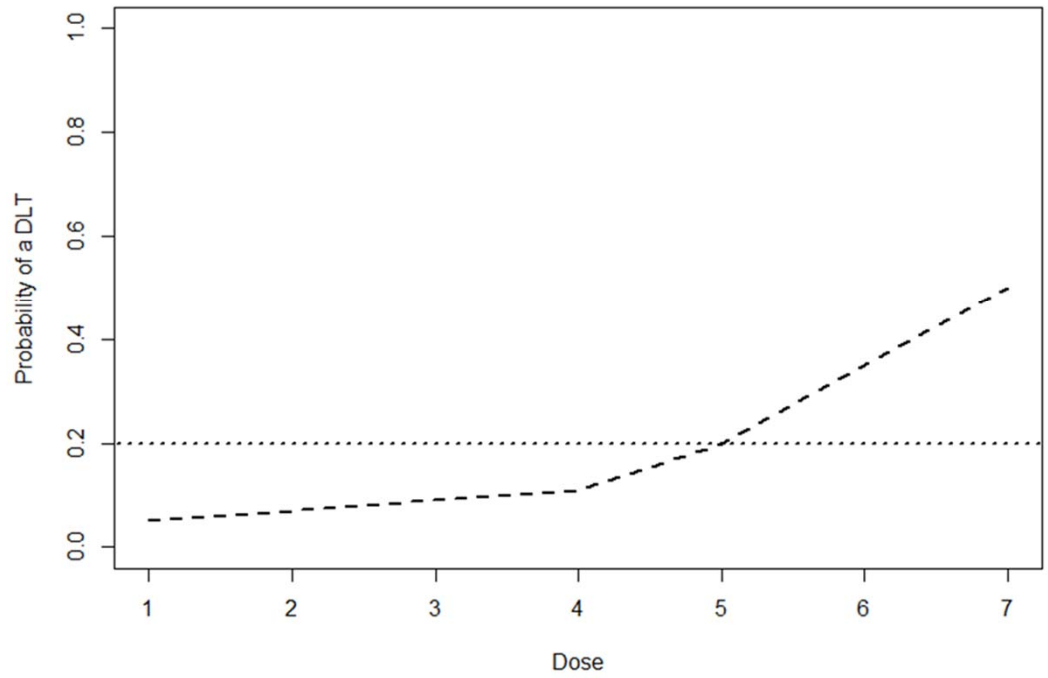

CRM Example: Recruit 3 Subjects and Update

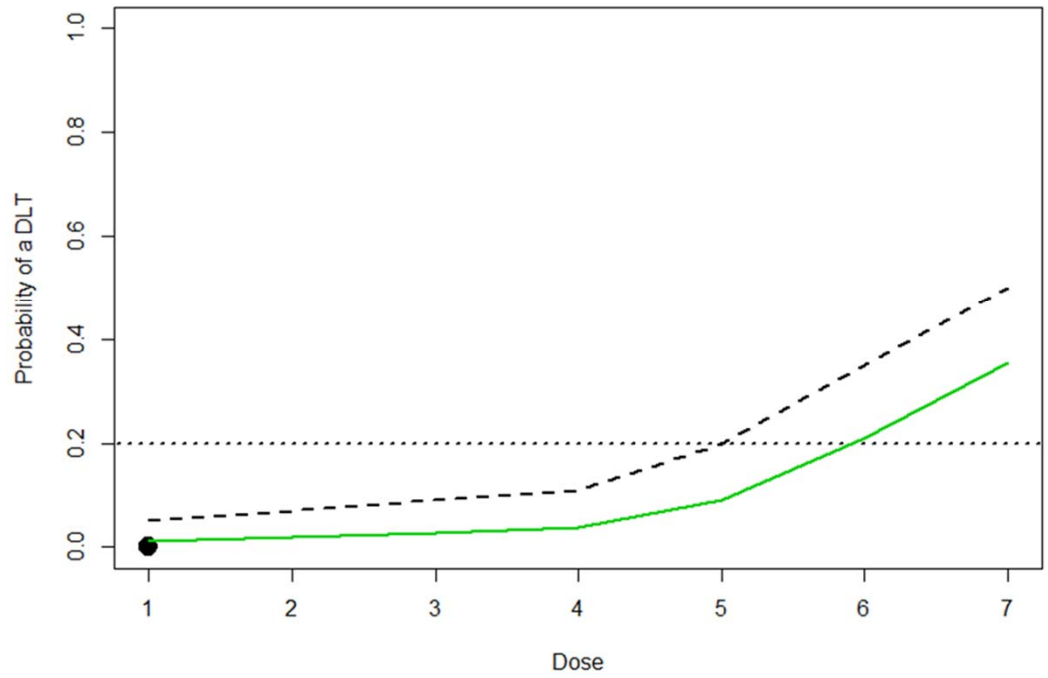

CRM Example: Current State

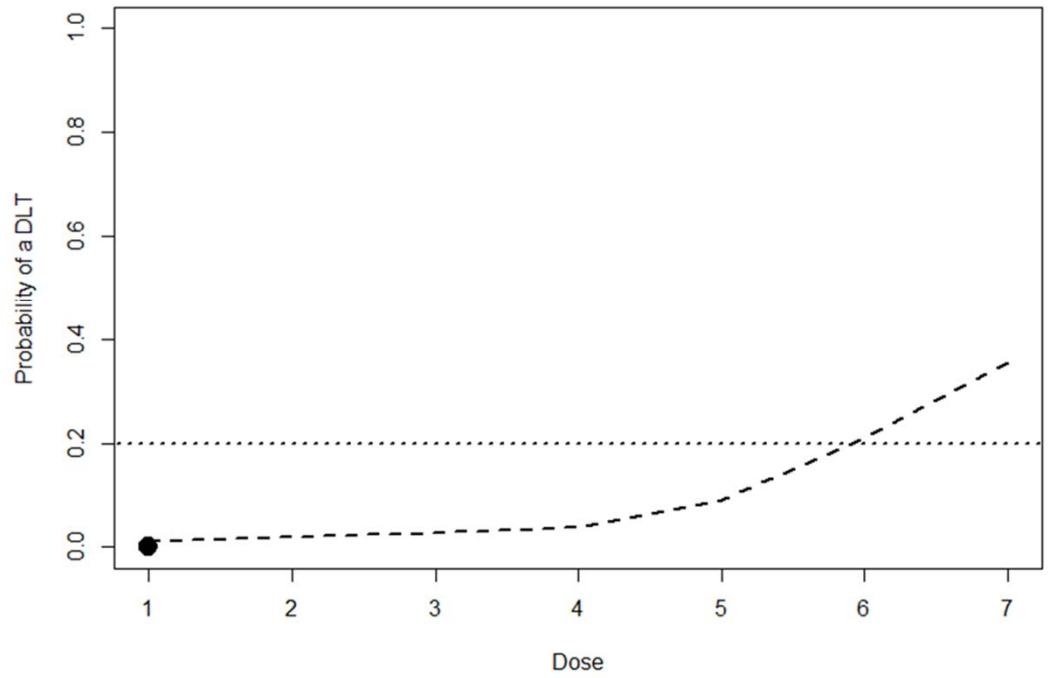

CRM Example: Recruit Another 3 Subjects and Update

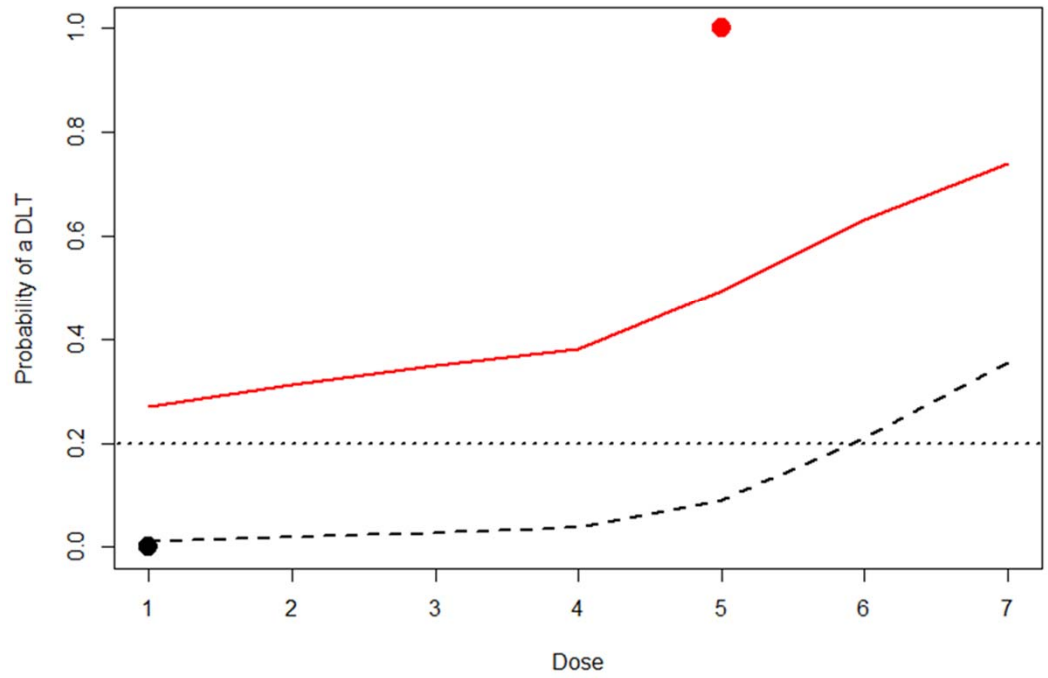

CRM Example: Current State

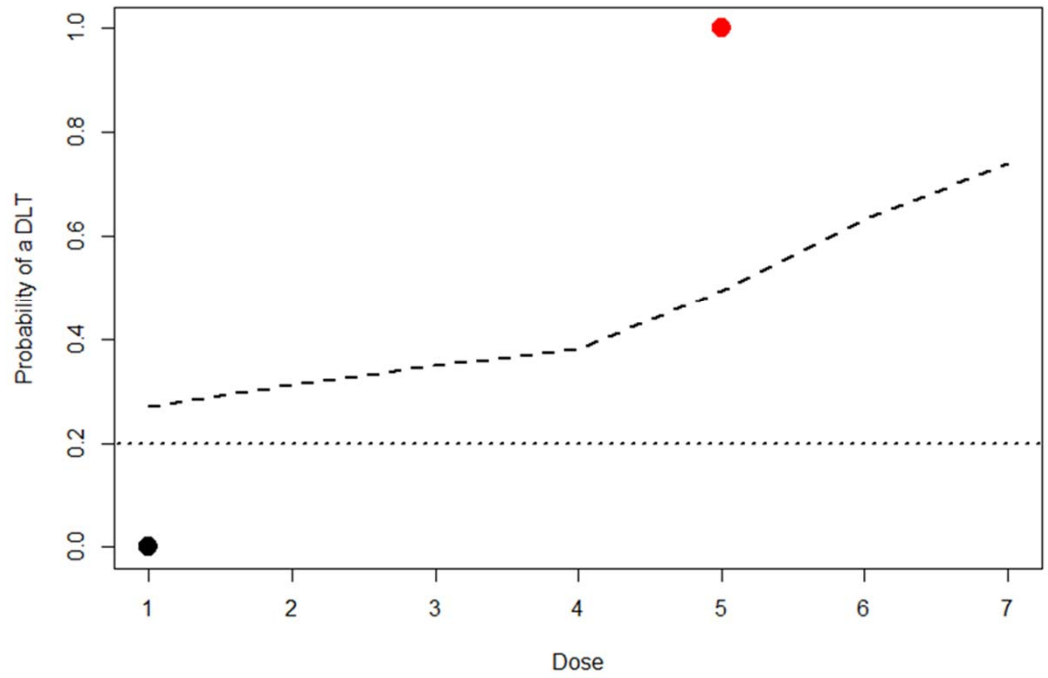

CRM Example: Recruit Another 3 Subjects and Update

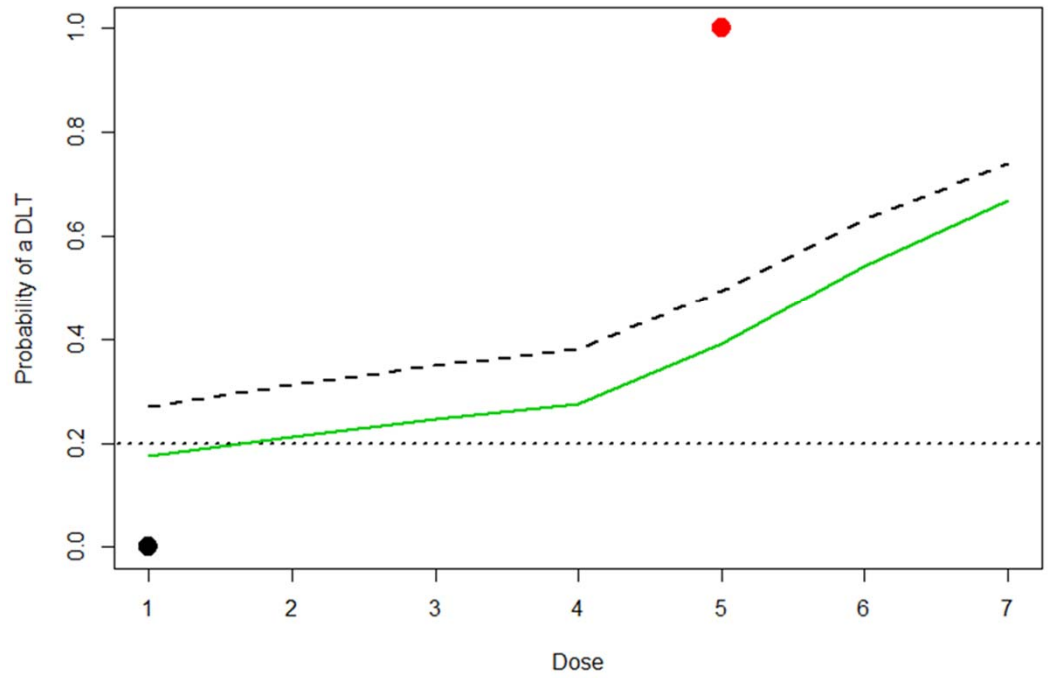

CRM Example: Continue Until You Reach a Stopping Criteria

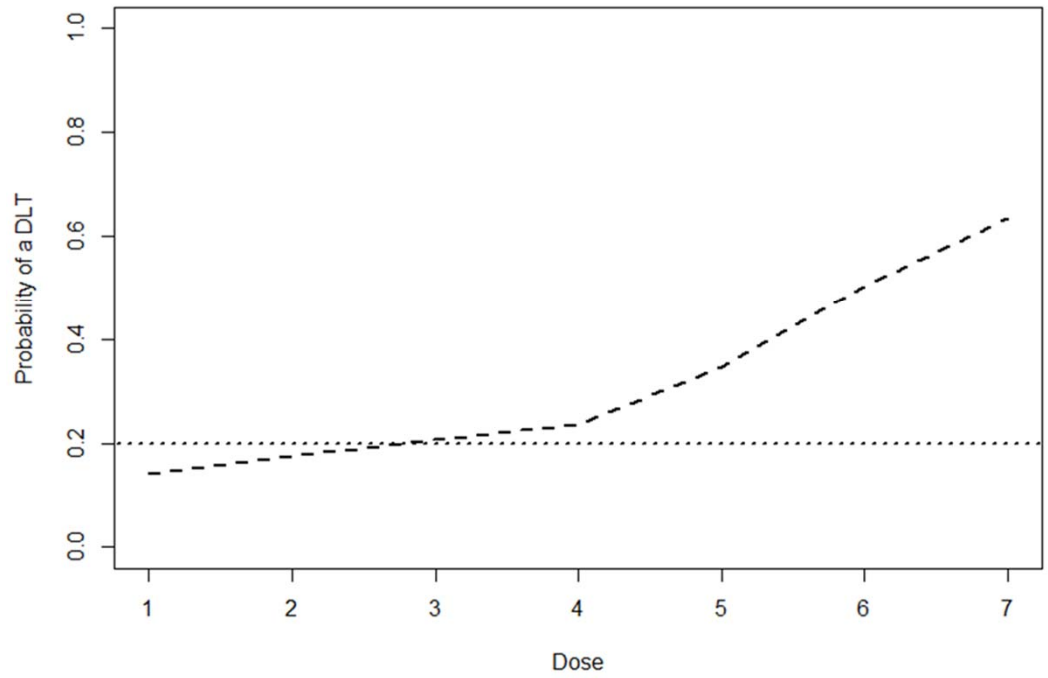

Questions?

References

- Doing Bayesian Data Analysis (2nd Edition) by John Kruschke