Tom Keelin
650.465.4800 (cell)
tomk@keelinreeds.com
www.metalogs.org

A Decision Analytic Framework for Bayesian Updating of Probability of Success in Clinical Trials

Probability of Success Interest Group

Society of Decision Professionals

January 17, 2024

Background (I)

- Metalog Bayesian inference was published as a preprint in 2021.
  - 1,359 downloads, 762 views
- Eric and Shaun inquired as to whether this method could be useful in calculating
  - probability of success (POS) for clinical trials
  - Bayesian updating of POS when new data becomes available

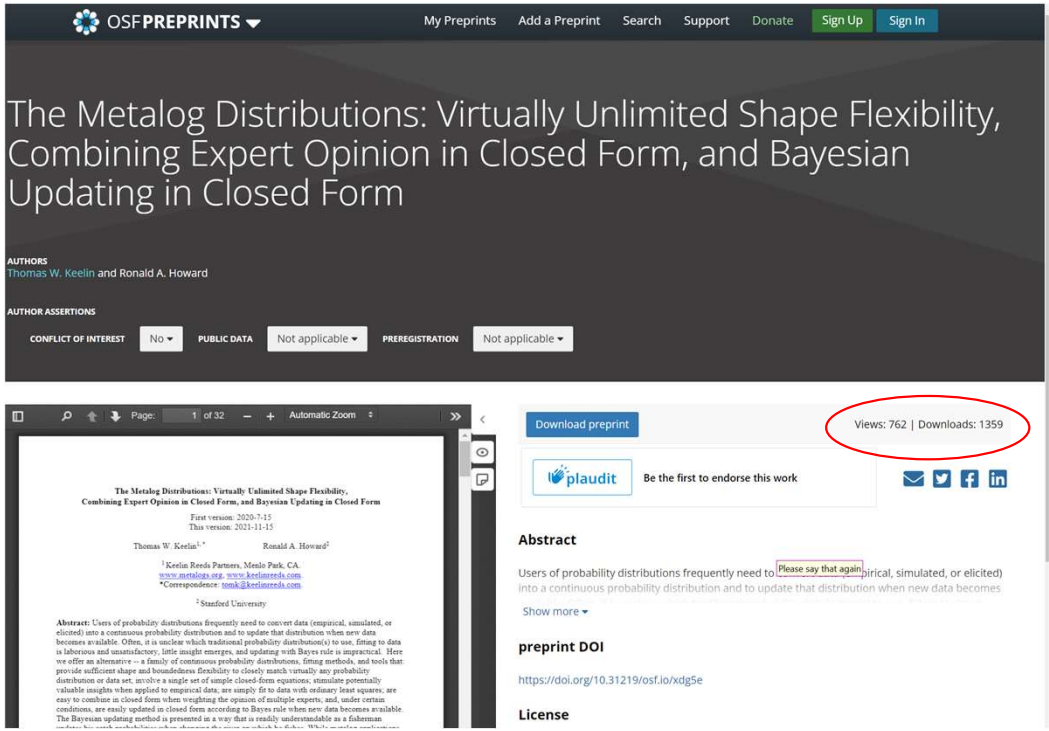

Google: Keelin and Howard OSF
Link: https://osf.io/preprints/osf/xdg5e

Example: New Weight-Loss Drug

Phase 1 Data

WTL %
1.60
3.21
8.11
1.01
4.02
2.67
3.62
3.45
5.64
4.24
2.22
1.87
4.61
-0.81
2.27
0.09

?	Control
3.1	Average

- In Phase 1 clinical trials, 16 patients lost an average 3.1% of their body weight after 6 months with minimal side effects
- Phase 2 trials are underway
- A possible phase 3 trial is being planned
  - 100 treated patients + 100 control patients
  - "Average weight loss" = Treated patient average weight loss – control group average
  - Success: Average weight loss >= 3% (success criterion)

Note: For now, we avoid hypothesis tests and other statistical success criteria.

Questions Posed by Shaun and Eric

- How can we calculate Phase 3 Probability of Success based on Phase I state of information (Phase 1 SOI)?

- How can we conveniently update this probability based on Phase 2 data (Phase 2 SOI)?

- Under what conditions is this updating procedure Bayesian?

A natural starting point is to consider our distribution over weight loss for the next patient(s).

Phase 1 Data

WTL %
1.60
3.21
8.11
1.01
4.02
2.67
3.62
3.45
5.64
4.24
2.22
1.87
4.61
-0.81
2.27
0.09

Distribution Over Net Weight Loss for Next Patient(s) (Ph 1 SOI)

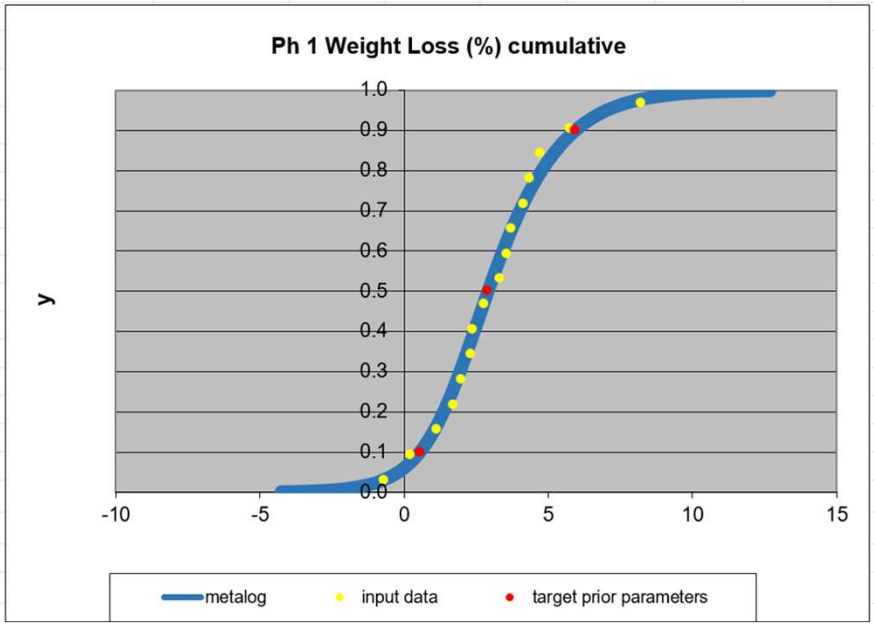

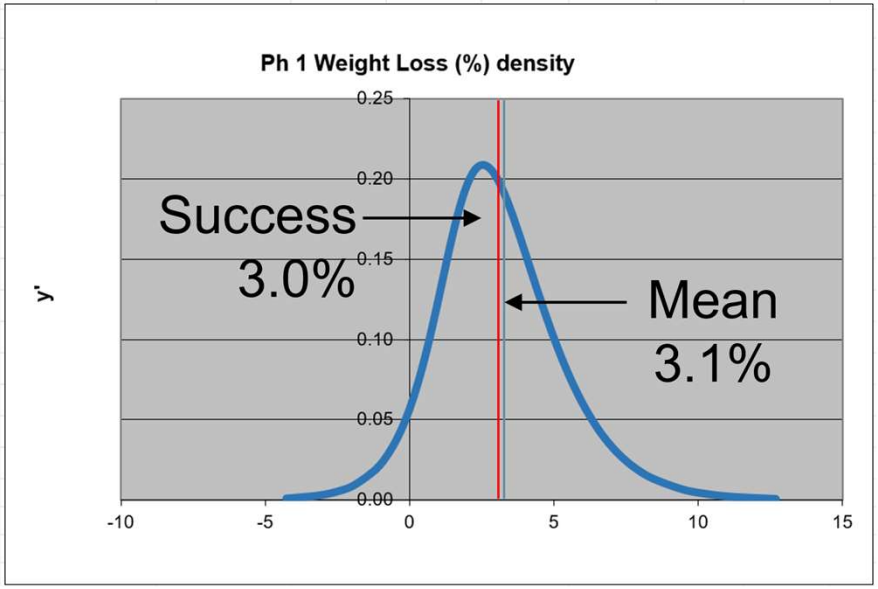

0 Control (for now)

3.1 Average

If this distribution were representative of the entire population, Phase 3 probability of success could be conveniently calculated by simulation.

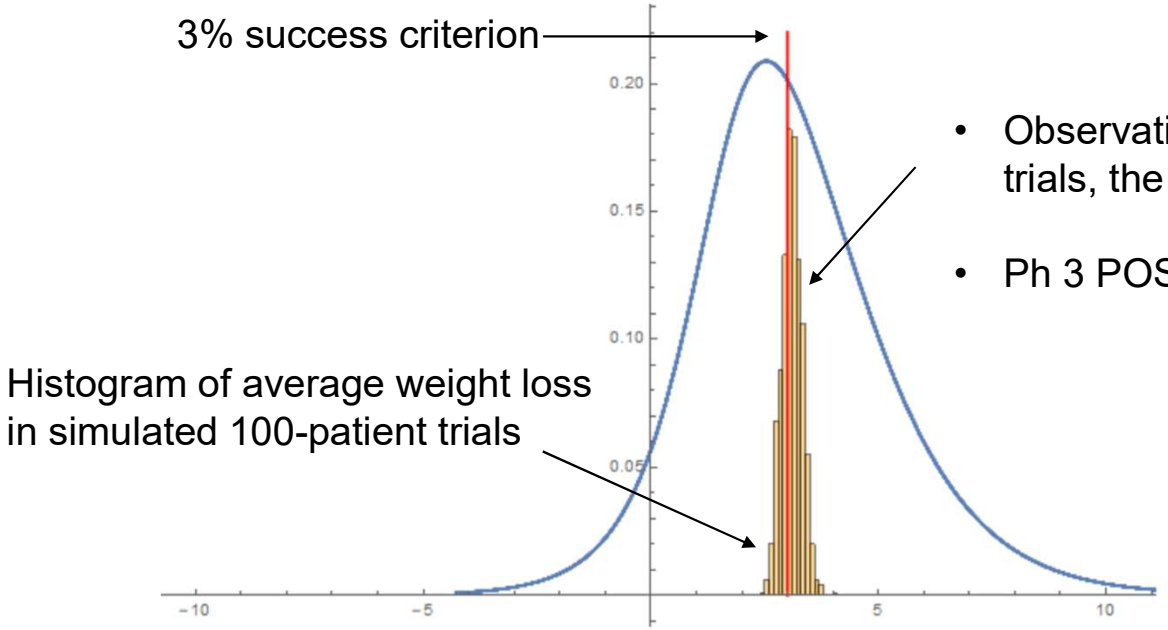

- Observation: in 68% of simulated 100-patient trials, the average weight loss > 3%
- Ph 3 POS (Ph 1 SOI) = 68%

If this distribution were representative of the entire population, increasing the number of patients in phase 3 would increasingly guarantee success.

100-patient trial

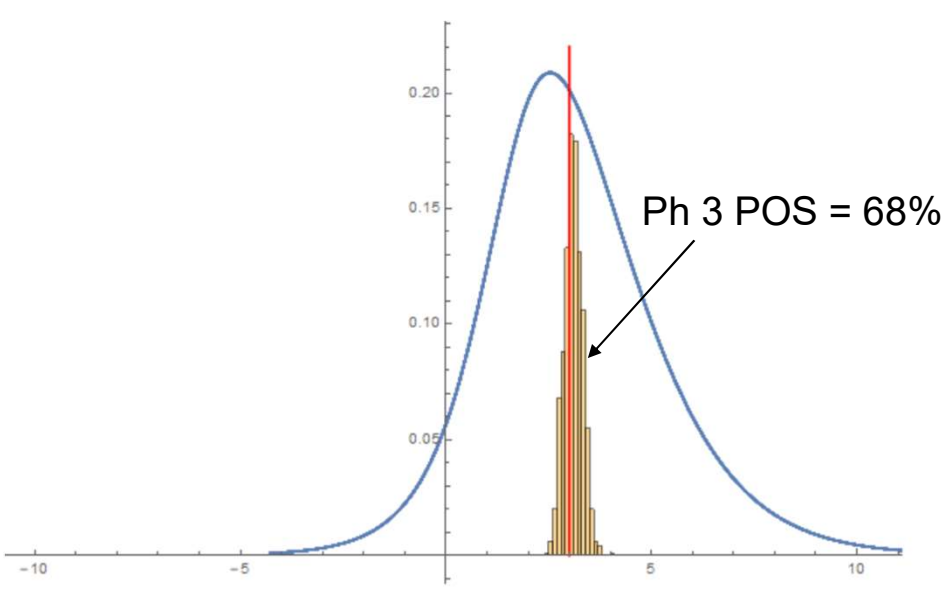

1000-patient trial

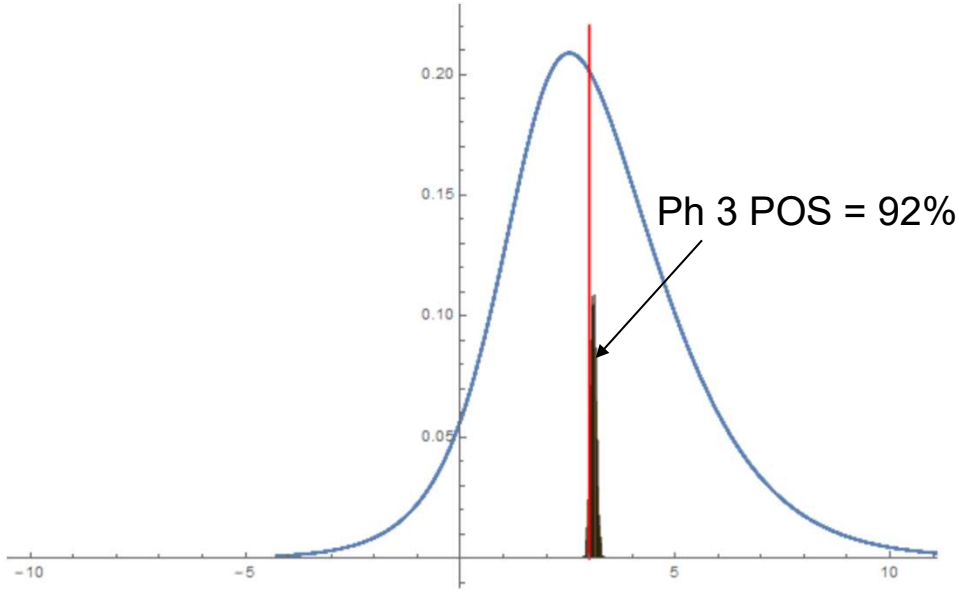

Similarly, if 0.5% placebo effect is assumed, POS drops to 3% and increasing the number of patients increasingly guarantees failure.

Phase 1 Data

WTL %
1.60
3.21
8.11
1.01
4.02
2.67
3.62
3.45
5.64
4.24
2.22
1.87
4.61
-0.81
2.27
0.09

0.5 Control

2.6 Average

100-patient trial

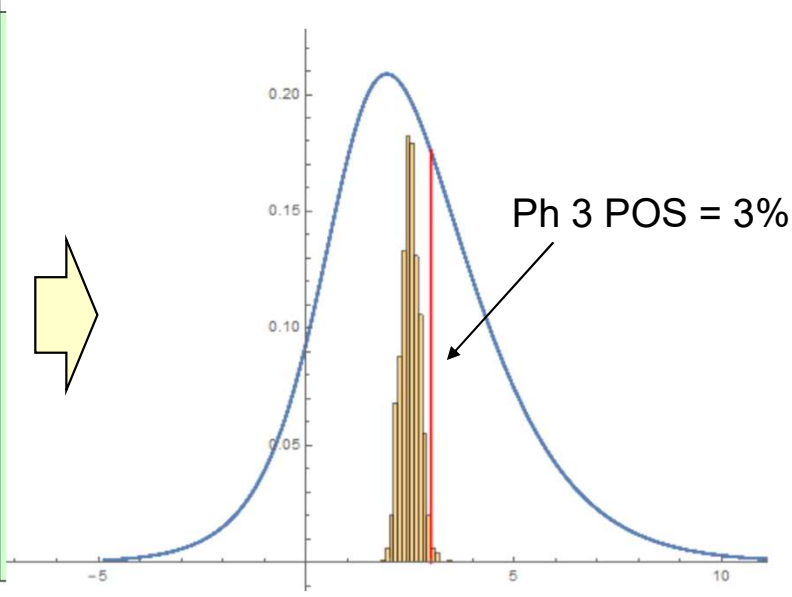

1000-patient trial

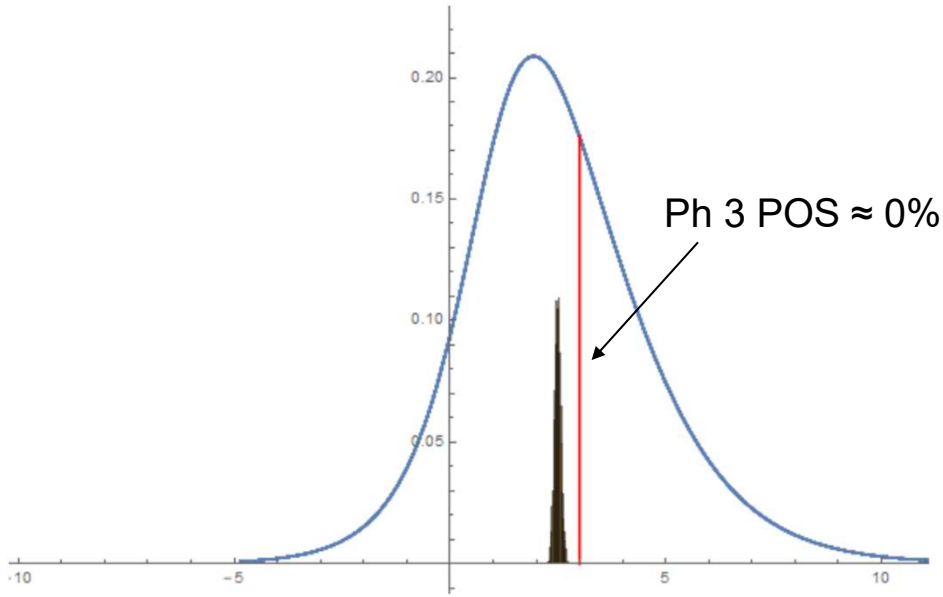

(assuming this distribution is representative of the entire population)

Based only on Ph 1 SOI, many different distributions over the entire population are possible.

Possible Entire-Population Weight Loss Distributions for Next Patient (Ph 1 SOI)

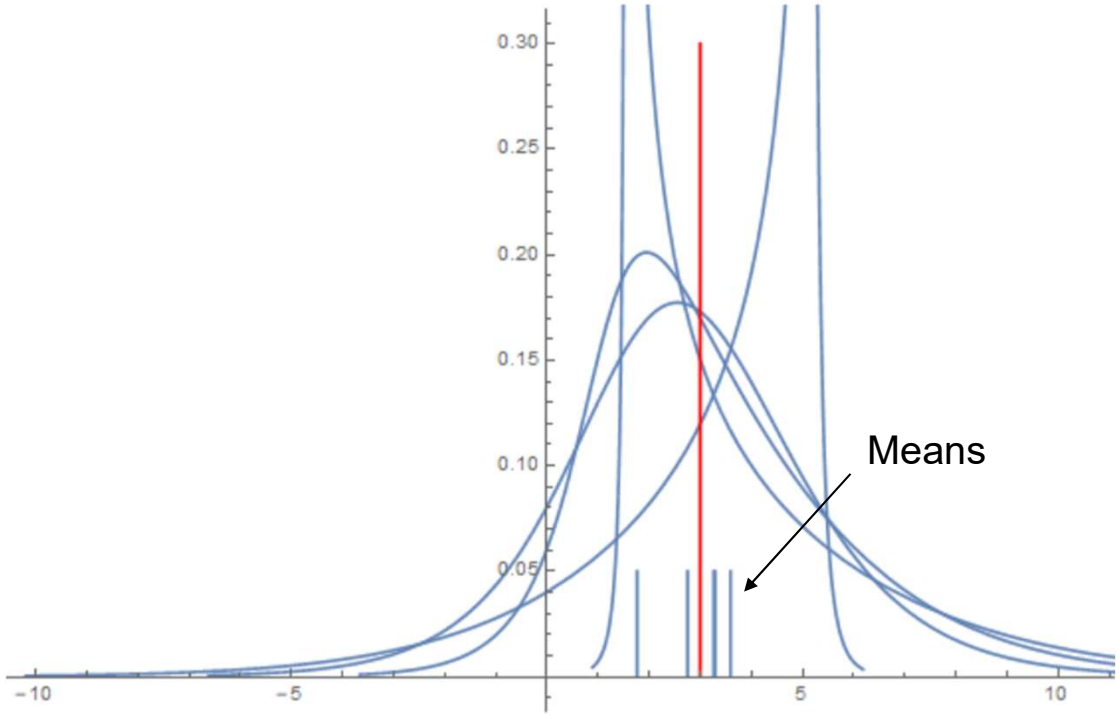

How might we generate a probability distribution over entire-population probability distributions (based on a SOI*)?

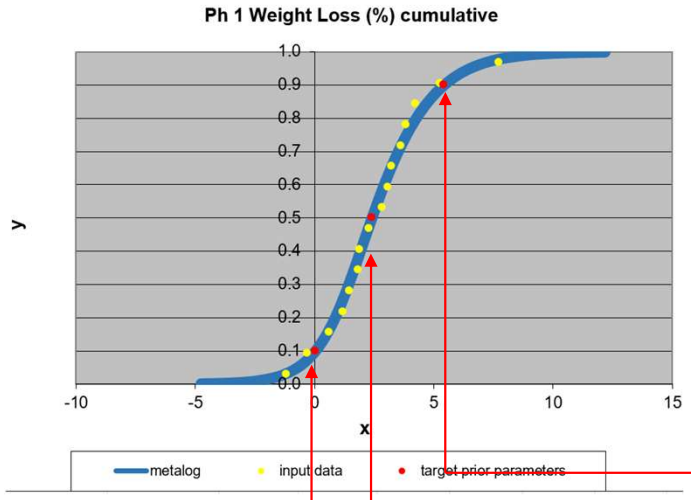

Probability distribution over the *parameters* of the entire-population weight-loss distribution

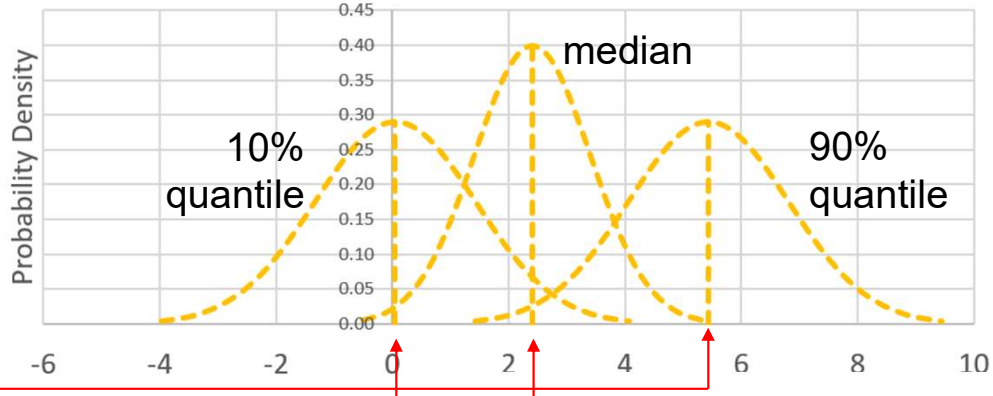

5.4

2.4

0.0

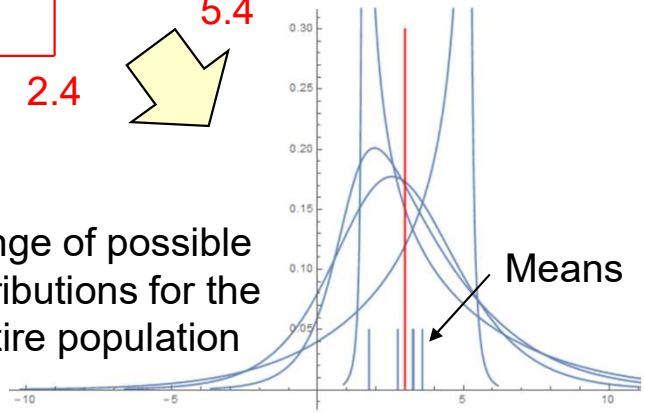

Range of possible distributions for the entire population

Means

* Henceforth, 0.5% Phase 1 placebo effect assumed unless otherwise indicated.

Procedure for calculating POS for a given state of information.

POS calculation method

A. Generate a sample weight-loss distribution from the distribution over parameters of the entire-population distribution

B. Conditional on that distribution, generate an n-patient sample of weight loss

C. Apply success criteria* to that sample (e.g. Success = Average(B) >= 3% Success Criterion)

D. Do A-C N times

E. POS = #Successes/N

*Success criteria can be *any* function of the patient data (e.g. statistical significance criteria or “more than 50% of patients lost more than 4% of body weight in excess of placebo”.)

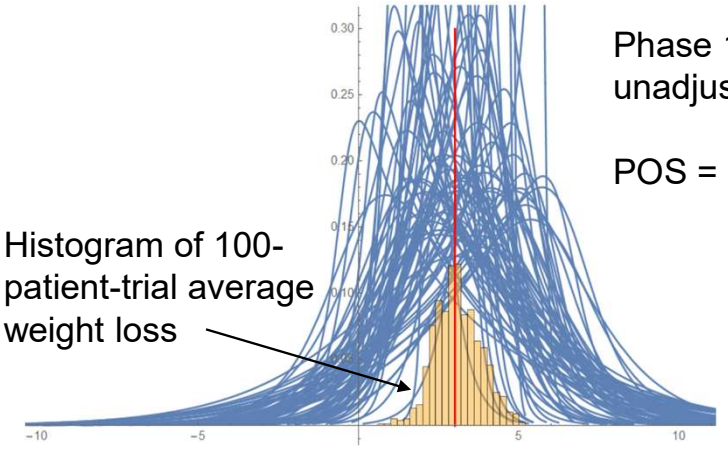

Phase 1 SOI unadjusted

POS = 52%

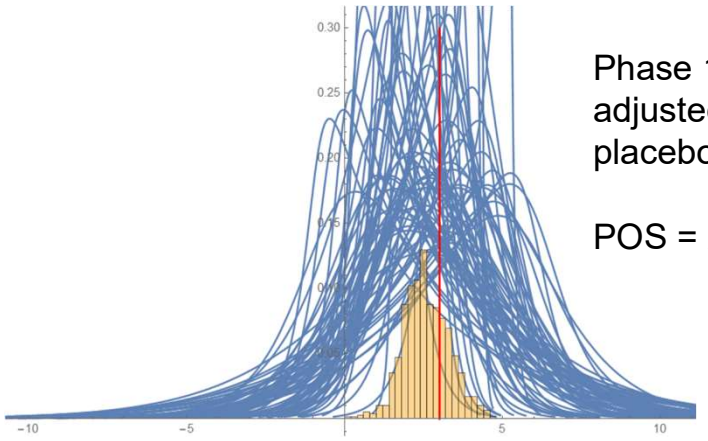

Phase 1 SOI adjusted for 0.5% placebo effect

POS = 28%

Would sampling clinical trials from the “average” entire-population probability distribution be a reasonable shortcut?

- Choice of averaging method is arbitrary
- Any such average will have a mean
- POS will go to either 0 or 1 as patient sample size increases, which it should not

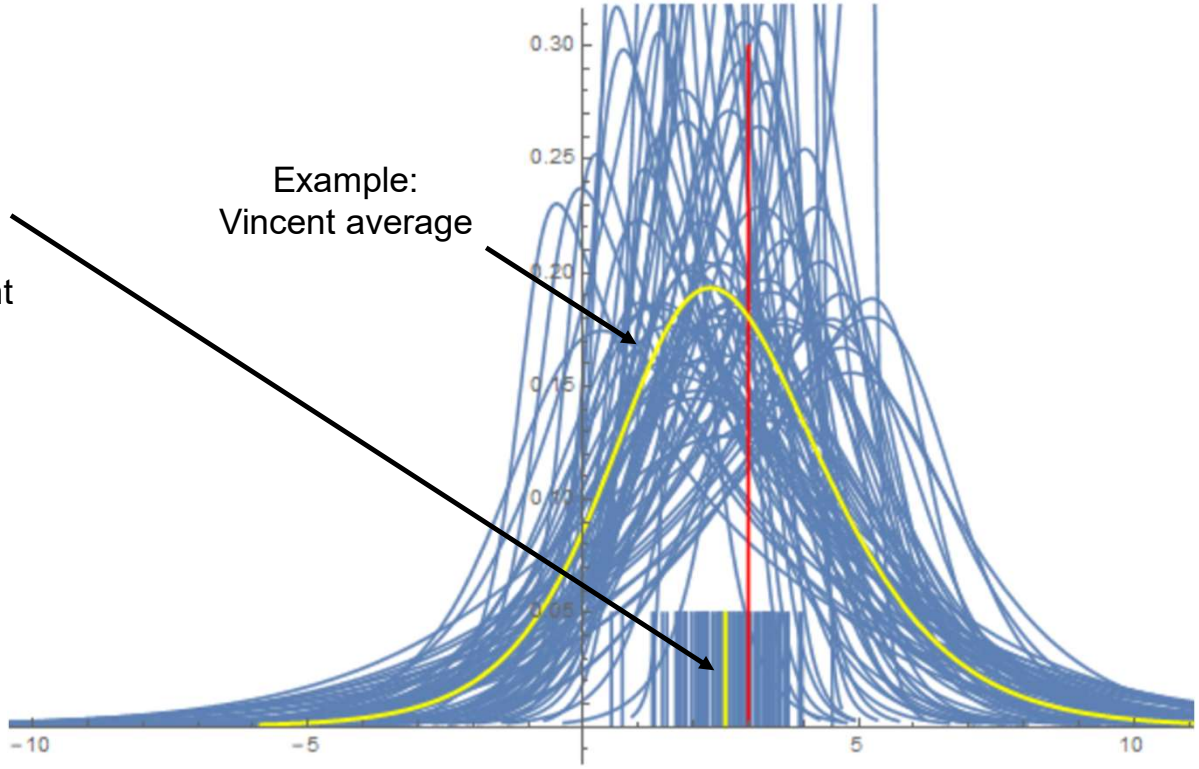

In contrast, creating a distribution of the *means* of the entire-population distributions can provide a good proxy for POS given a state of information.

A. Generate a sample from the distribution over parameters of the entire-population distribution

B. Calculate the mean of that distribution

C. Apply success criteria (e.g. Success = mean B >= 3% Success Criterion)

D. Do A-C N times

E. POS = #Successes/N

Phase 1 SOI adjusted for 0.5% placebo effect

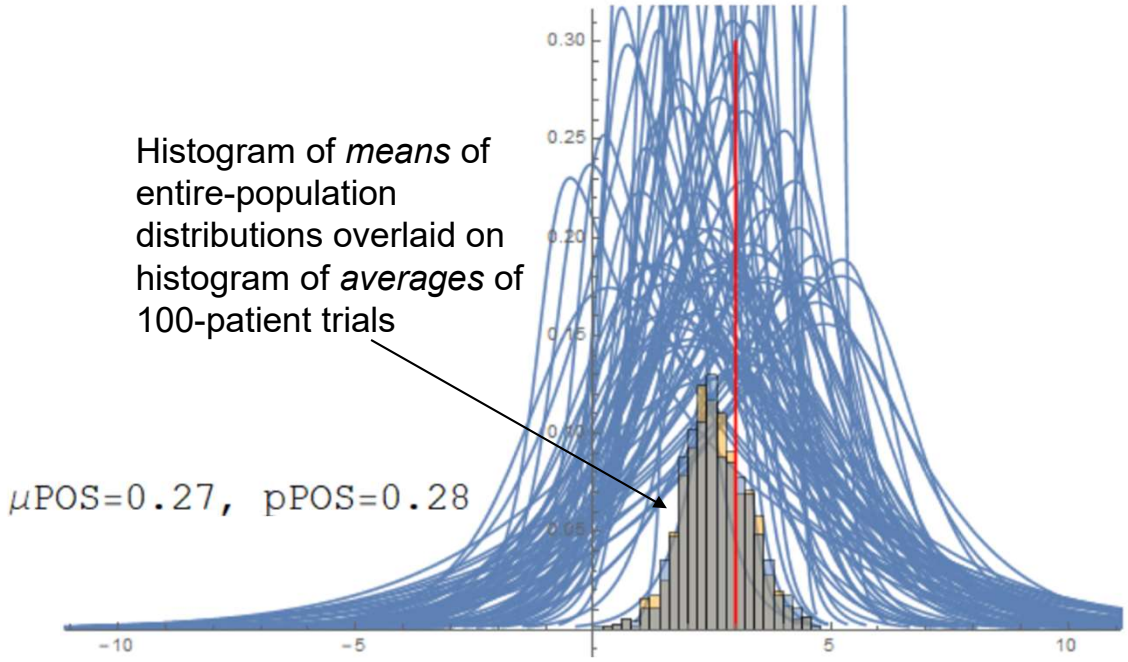

Note: The larger the number of patients in Phase 3, the more reliable is this proxy.

Questions Posed by Shaun and Eric

- How can we calculate Phase 3 Probability of Success based on Phase I state of information (Phase 1 SOI)?

- How can we conveniently update this probability based on Phase 2 data (Phase 2 SOI)?

- Under what conditions is this updating procedure Bayesian?

Shaun and Eric generated two cases of Phase 2 efficacy data for consideration. No safety concerns were observed in either case.

Favorable Case

Phase 2 Data

WTL %
3.66
6.13
4.44
3.89
8.00
2.68
1.60
1.34
5.33
4.52
5.03
3.45
4.84
2.71
-3.61
7.36
5.27
5.13
3.72

0.25	Control
3.73	Average

Unfavorable Case

Phase 2 Data

WTL %
5.64
1.41
4.52
5.46
4.72
5.10
5.24
5.80
2.62
3.49
3.12
5.96
7.50
-0.70
5.76
4.79
-0.02
2.51
2.48

1.15	Control
2.82	Average

New data enables us to update the distribution over the parameters and POS (I).

Favorable Case

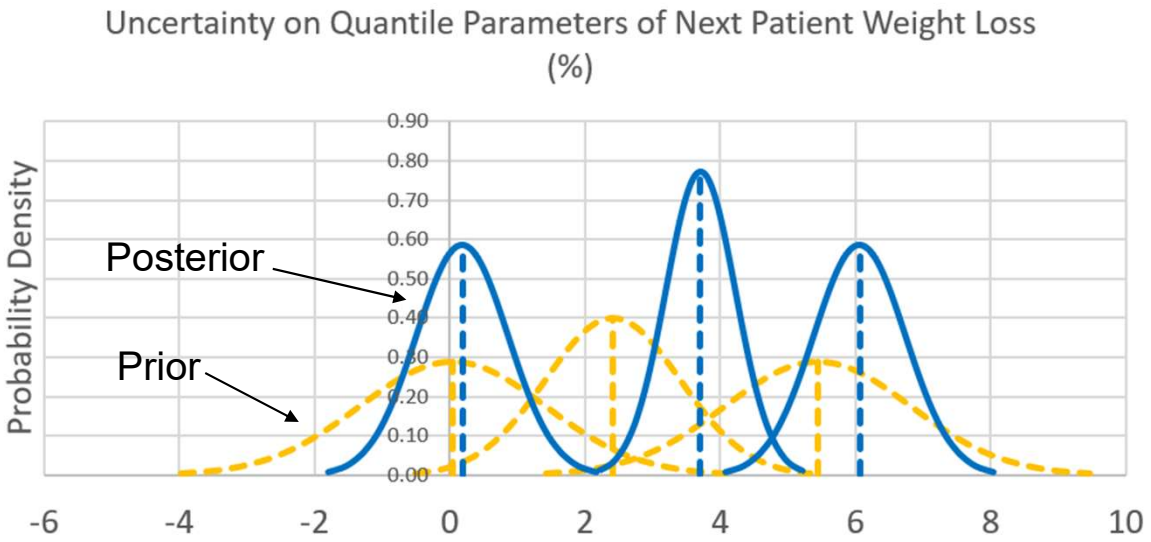

Phase 2 SOI

POS = 80%

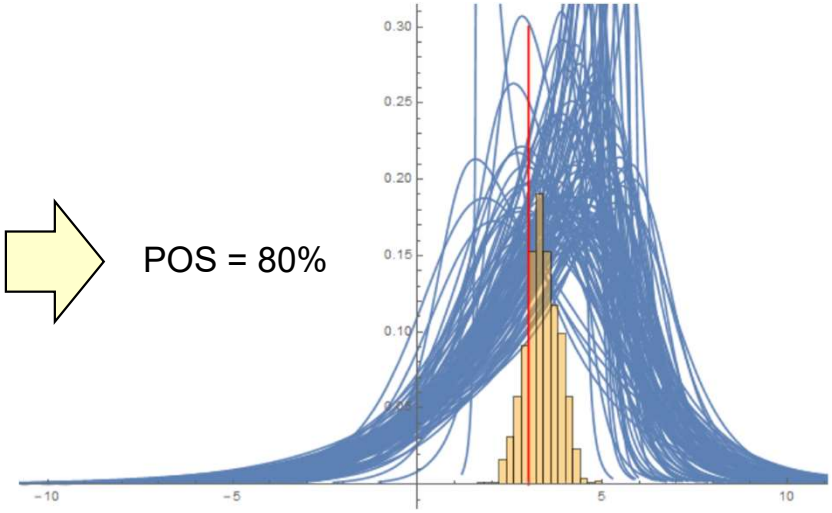

Updating the mean: An appealing method is to combine *representative* prior data with new data and define the updated mean as the metalog parameters (fit with least squares) of the combined data.

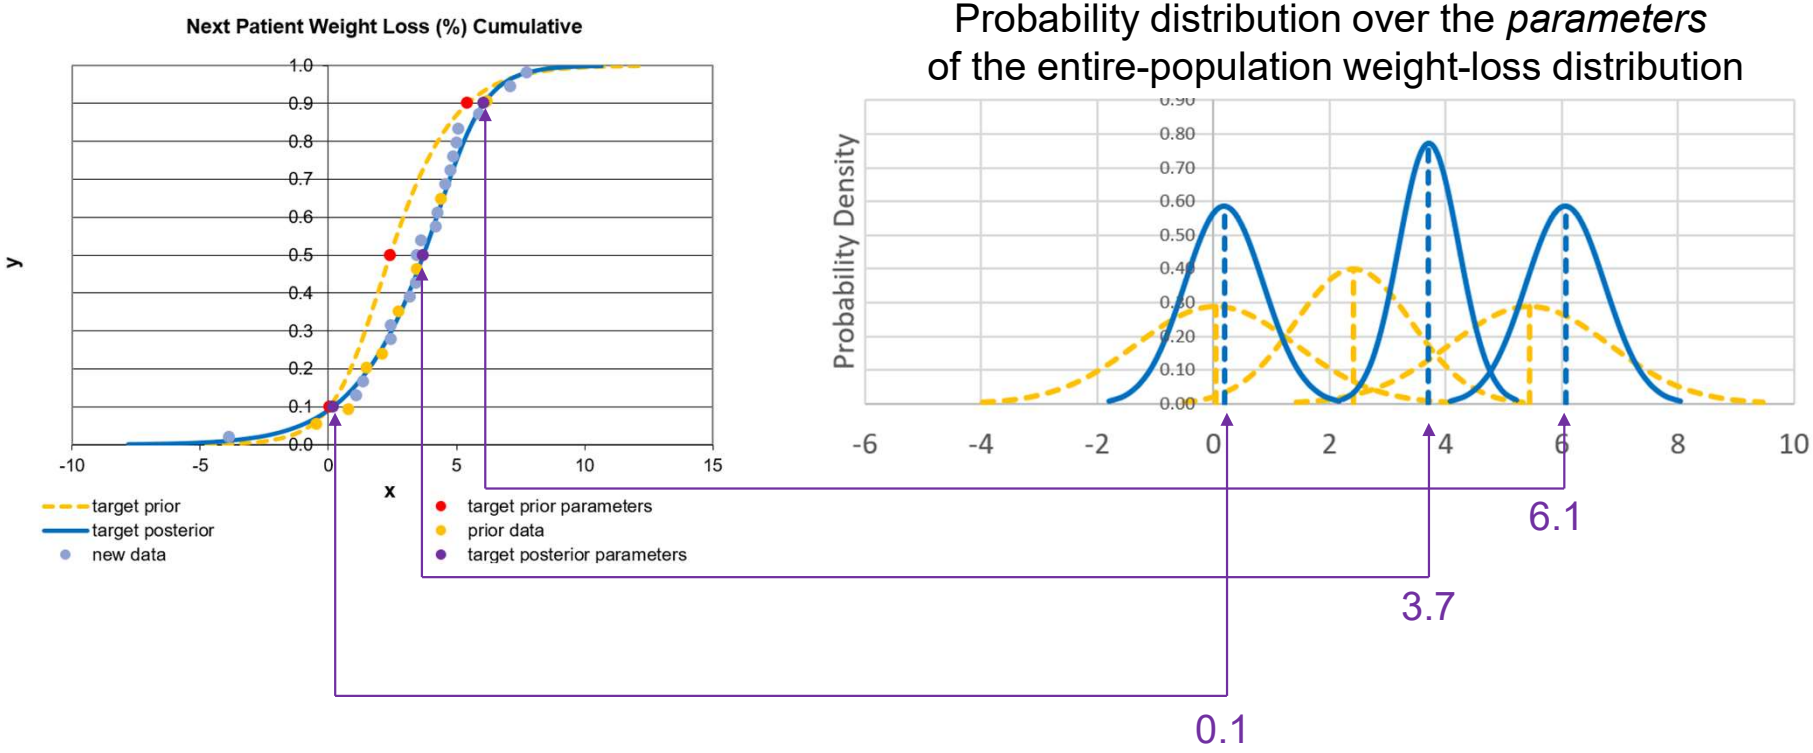

Updating the mean: What *representative* prior data would we use?

Phase 1 Data

WTL %
1.60
3.21
8.11
1.01
4.02
2.67
3.62
3.45
5.64
4.24
2.22
1.87
4.61
-0.81
2.27
0.09

0.5 Control

2.6 Average

Option 1: Use all prior data

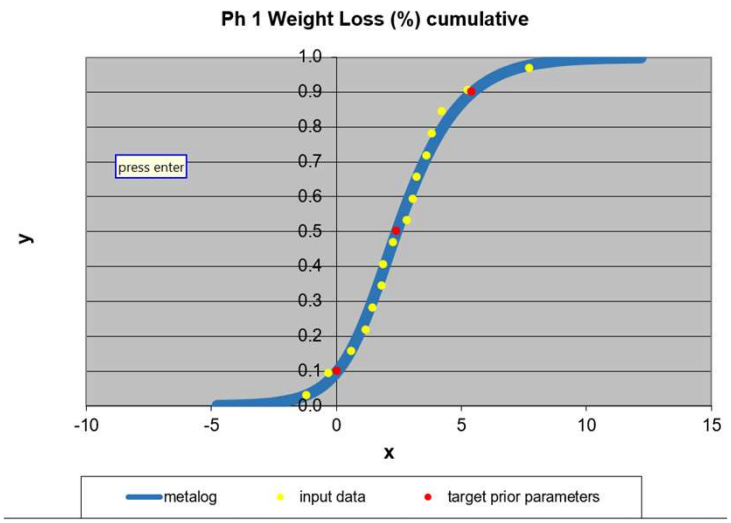

Option 2: Assess "equivalent sample size" n0. Use n0 data with the same shape and location

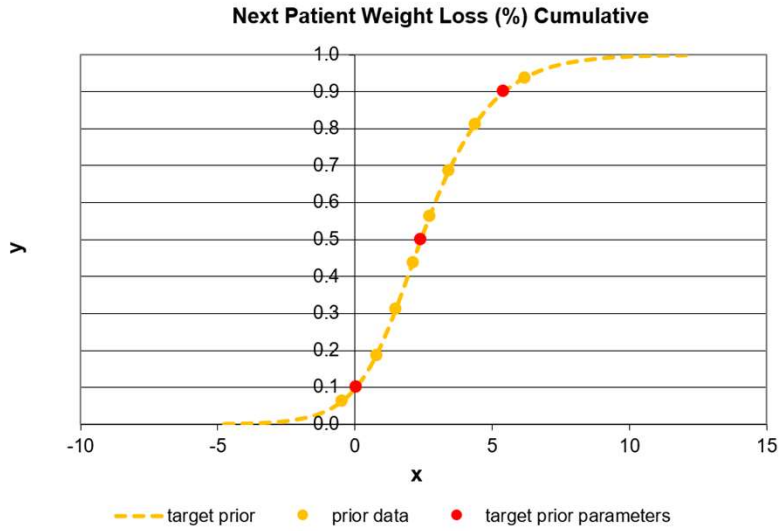

prior data
x
-0.47
0.78
1.50
2.11
2.72
3.43
4.38
6.18

Example: n0 = 8

(Option 2 with n0 = 8 is our base case.)

Specifying and updating the covariance matrix is even more convenient.

- Covariance matrix depends only on
  - initial width σ
  - total number of data n

σ: a scalar constant that determines width of the prior over parameters

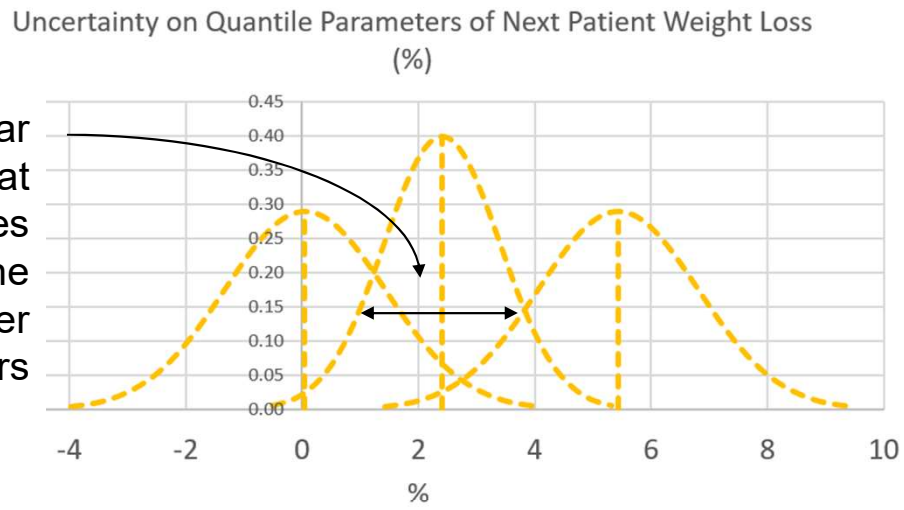

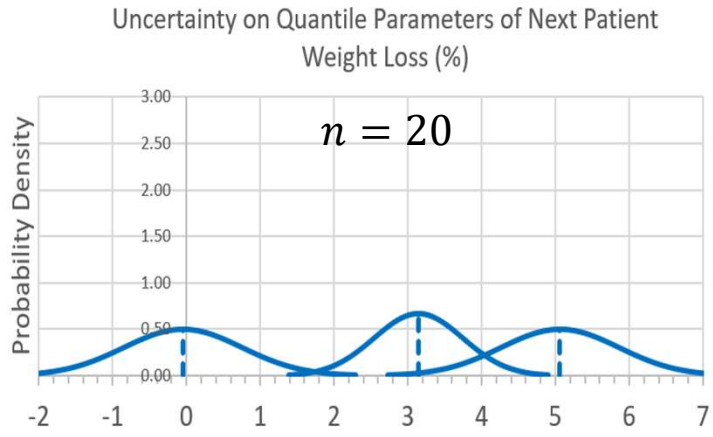

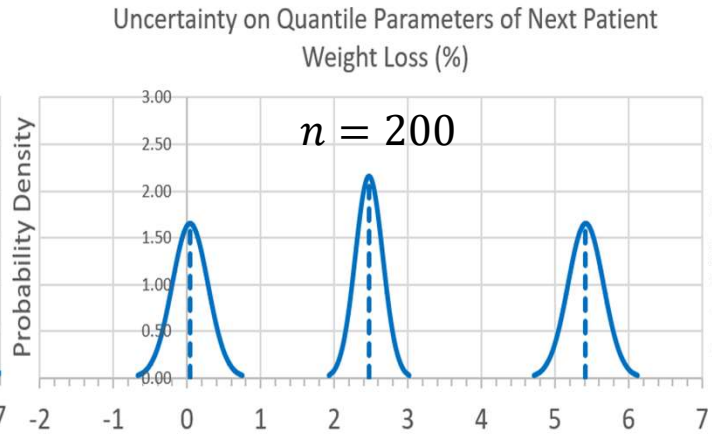

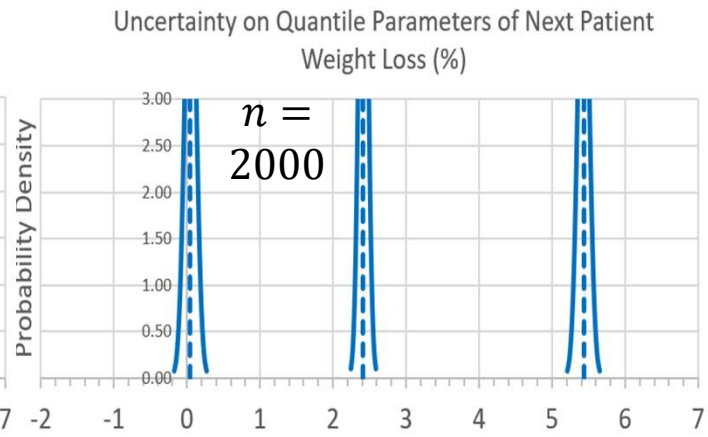

Note: variances and covariances shrink in proportion to $1/n$

In summary, based on these methods, the Phase 3 POS is ~80% given favorable Ph Phase 2 data.

Favorable Case

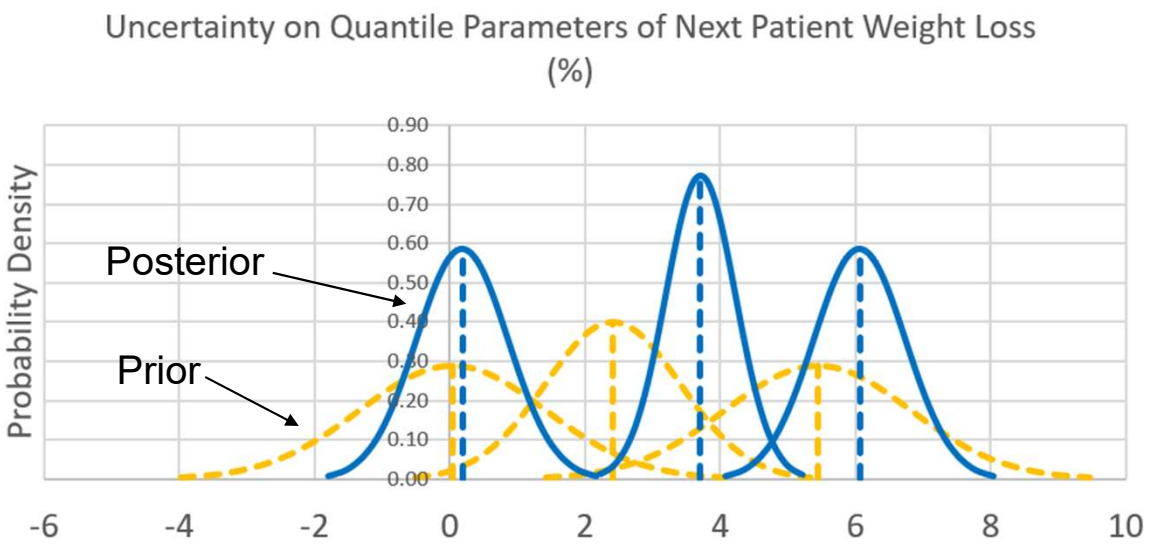

Phase 2 SOI

POS = 80%

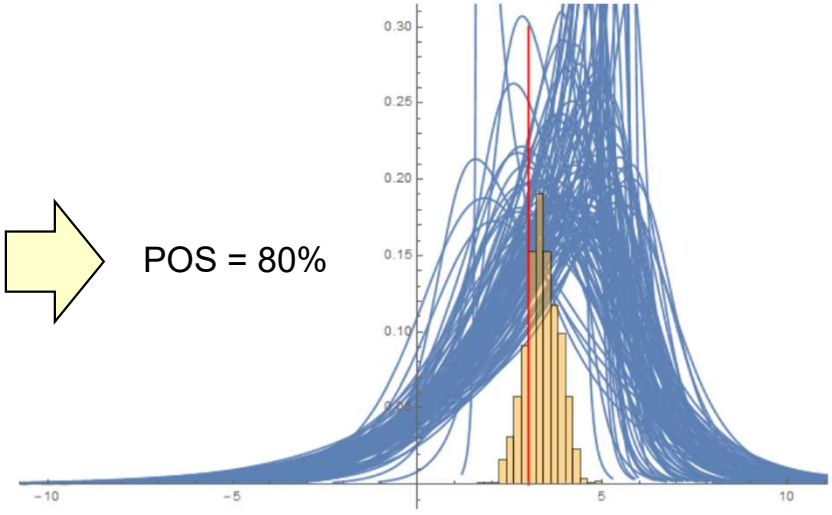

Sensitivity analyses show the effect of a range of parameter settings.

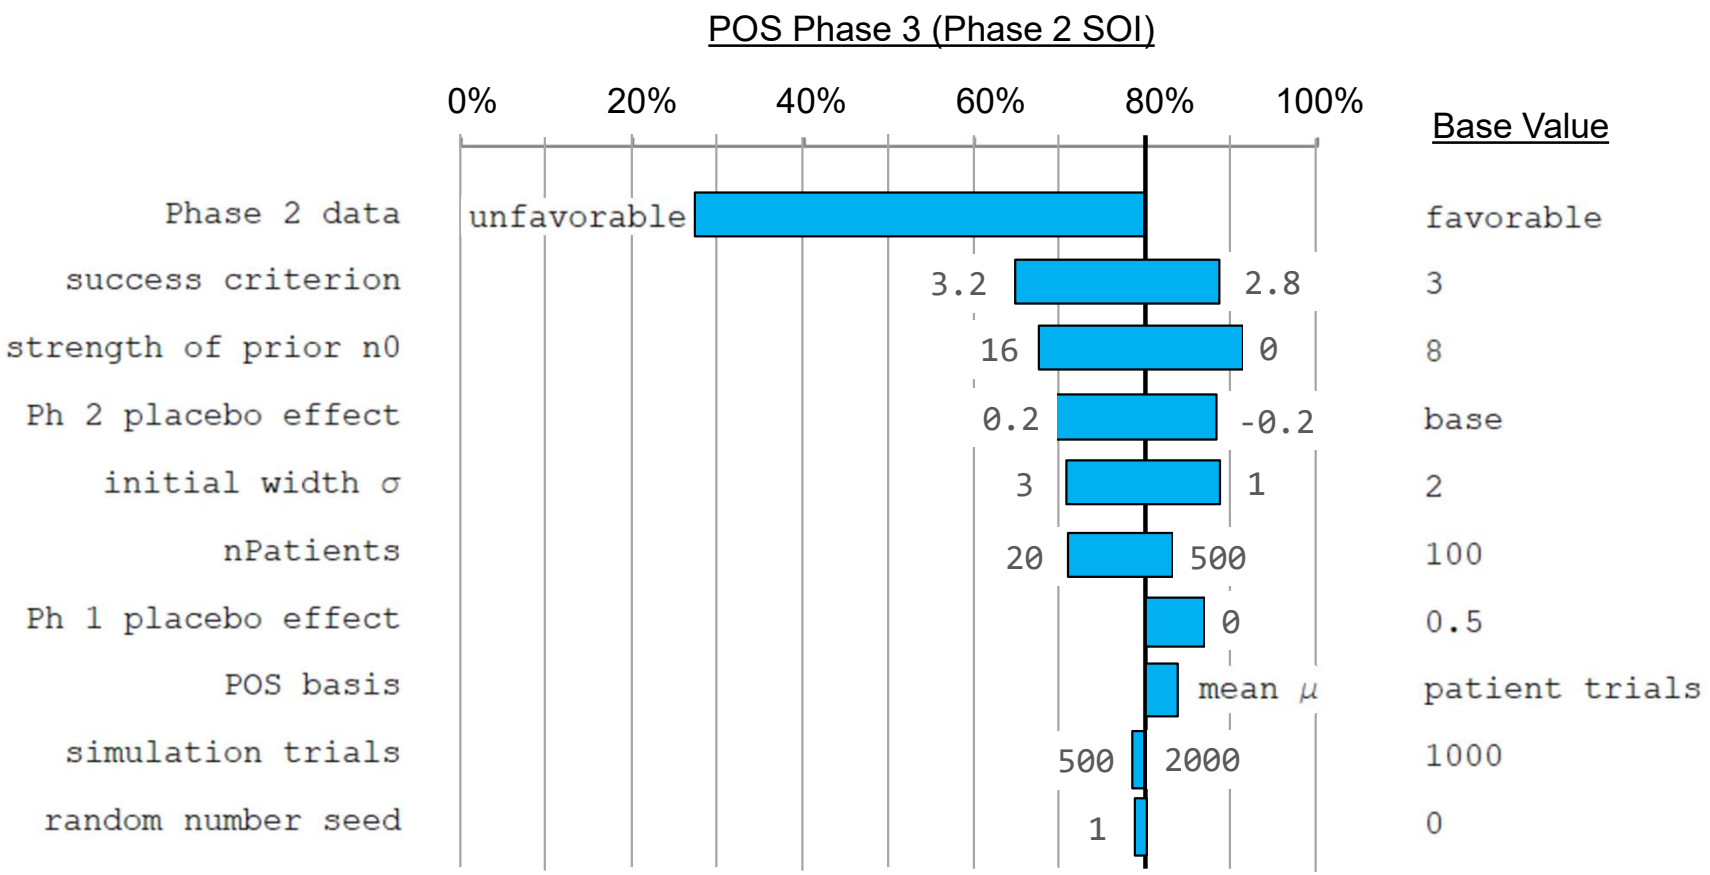

Questions Posed by Shaun and Eric

- How can we calculate Phase 3 Probability of Success based on Phase I state of information (Phase 1 SOI)?

- How can we conveniently update this probability based on Phase 2 data (Phase 2 SOI)?

- Under what conditions is this updating procedure Bayesian?

These simple, closed form updating methods are Bayesian under certain conditions.

1. Target variable distribution in a QPD

2. Distribution over parameters is multivariate normal (or multivariate Student t)

3. Sample "errors" are normally distributed with standard deviation σ and mean 0.

4. Prior data contains all relevant prior information.

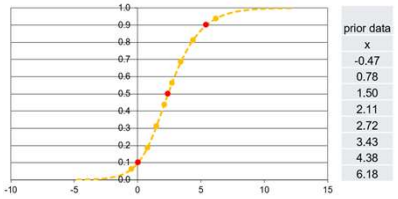

- Quantile function is linear in its parameters*.

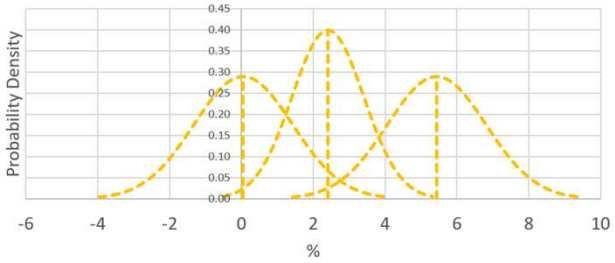

- Likelihood function is a product of normal distributions

- Distribution over parameters is a conjugate prior
  - σ certain -> multivariate normal
  - σ uncertain -> multivariate Student t

- Bayesian updating equations simplify to the above procedure.

Caveat: "All models are wrong, but some are useful" (George Box): A small fraction of samples from the distribution over parameters is infeasible. We discard those.

*Keelin and Powley, Quantile-Parameterized Distributions, Decision Analysis, 2011.

Questions Posed by Shaun and Eric

- How can we calculate Phase 3 Probability of Success based on Phase I state of information (Phase 1 SOI)?
- How can we conveniently update this probability based on Phase 2 data (Phase 2 SOI)?
- Under what conditions is this updating procedure Bayesian?

Thank you!