Efficient Simulation and Comparison of Adaptive Clinical Trial Designs

Peter Jacko*

PSI, London
4 Jun 2019

*Dept. of Management Science, Lancaster University, UK

This Talk

- Two treatments
- Binary endpoints
- "Immediate" responses
- Trial of n patients
- Response-adaptive clinical trial design
  - ▷ is a (possibly randomized) allocation strategy that specifies which treatment the i-th patient receives
  - ▷ can be frequentist or Bayesian
  - ▷ includes designs with fixed randomization

Randomised Controlled Trial

- The gold standard design: randomised controlled trial
  - 50% vs 50% fixed equal randomisation
  - avoids many types of biases
  - in use since 1948 (advocated since Hill 1937)
- Its main goal is to learn about treatment effectiveness with a view to prioritising future outside patients
  - maximises power of a treatment effect difference
  - if approved, future patients are, say, 95% confident that the novel treatment is better than the control
- A half of trial patients gets the inferior treatment

Randomised Controlled Trial

- Advantages of randomised controlled trial
  - ▷ best possible for estimating the treatment effect
  - ▷ "easy" to understand by trial statisticians, physicians
    - basic statistical knowledge
  - ▷ "quick" to design
    - trial size formulae available
  - ▷ "straightforward" to implement
    - no computation needed during the trial
  - ▷ "easy" to understand by in-trial patients (?)
  - ▷ "easy" to interpret the results by regulators
  - ▷ "easy" to interpret the results by physicians (?)
  - ▷ "easy" to interpret the results by patients (?)

Randomised Controlled Trial

- Disadvantages of randomised controlled trial
  - ▷ cost: 20% error of not approving a better treatment
    - — development and approval processes: $ billions
  - ▷ faith: once approved, no (simple) way to retract the treatment
    - — worse treatment approved by 5% chance
    - — unforeseen long-term secondary effects
  - ▷ feasibility: requires hundreds of patients for a trial
  - ▷ stationarity: approval process takes years
    - — inappropriate for new diseases and epidemics
  - ▷ ethics: patients join a trial expecting to get a possibly better (unapproved yet) treatment

Randomised Controlled Trial

- Statistical testing based on randomised equal allocation is a widespread state-of-the-art approach in the design of experiments, under different names:
  - randomised controlled trial in biostatistics
  - between-group design in social sciences
  - A/B testing in Internet marketing
- In other fields however the disadvantages are usually less severe

Bayesian Decision-Theoretic Trial

"...there can be no objection to the use of data, however meagre, as a guide to action required before more can be collected ... Indeed, the fact that such objection can never be eliminated entirely—no matter how great the number of observations—suggested the possible value of seeking other modes of operation than that of taking a large number of observations before analysis or any attempt to direct our course..."

- Proposed in Thompson 1933 (pre-dates Hill 1937)

Health Benefit Approach

- The goal is to provide higher health benefit to both in-trial patients and after-trial patients
  - ▷ healing patients is the ultimate goal of new treatment development
- As opposed to RCT's learning goal of reliable treatment effect estimation
  - ▷ the estimate will never be accurate
  - ▷ the treatment will be replaced by a new one in future

Bayesian Decision-Theoretic Trial

- Bayesian decision-theoretic model
  - ▷ optimally solving learning/healing trade-off
  - ▷ both learning and healing takes place during the trial
- It is done by deciding the allocation, i.e., the randomisation probabilities for every patient
  - ▷ Bayesian response-adaptive: decisions based on the responses accumulated so far

Bayesian Decision-Theoretic Trial

- In theory, can be solved to optimality by decision theory
  - it is often believed tractable only for small trials
- In general known as the multi-armed bandit problem
- Milestones IMHO
  - Thompson (Biometrika 1933)
  - Glazebrook (Biometrika 1978)
  - Gittins & Jones (Biometrika 1979)
  - Armitage (ISR 1985)
  - Cheng, Su & Berry (Biometrika 2003)
  - Berry (Nature 2006), Cheng & Berry (Biometrika 2007)
  - Villar, Bowden & Wason (Statistical Science 2015)

Multi-Armed Bandit Problem

Multi-armed Bandit Problem

- Optimally solving learning/earning trade-off
- Studied in other scientific disciplines including probability, statistics, operational research, economics, marketing, machine learning, computer simulation, computer science, and communications engineering
- Many formulations: i.i.d., Markovian, Bayesian, etc.
- Many extensions, mainly in machine learning
- Appropriate model for trials: finite horizon
  - ▷ after the end of the horizon we will not be able to influence the allocation

Bayesian Bernoulli Bandit Model

- Finite horizon: n sequentially arriving patients
- Two-armed: treatment C or D for each patient
- Binary endpoints: success (1) or failure (0)
- Let X_i and Y_i denote patient i's response from treatment C and D respectively (for $i = 1, ..., n$). Then,

$$X_i \sim \text{Bernoulli}(1,\ \theta_C) \ \text{ and } \ Y_i \sim \text{Bernoulli}(1,\ \theta_D),$$

where θ_C and θ_D are the unknown success probabilities of treatments C and D respectively

Bayesian Approach

- Beliefs $\widehat{\theta}_C$ and $\widehat{\theta}_D$ to be updated over the trial
- Prior Distribution: $\widehat{\theta}_C \sim \text{Beta}(a,\ b)$, $\widehat{\theta}_D \sim \text{Beta}(c,\ d)$ where we take $a = b = c = d = 1$ (uninformative)
- Posterior Distribution: After observing i (j) successes (failures) on treatment C, and k (l) successes (failures) on treatment D, the posterior distribution is represented by another Beta distribution (by conjugacy)

$$\widehat{\theta}_C \sim \text{Beta}(a+i, b+j),\ \widehat{\theta}_D \sim \text{Beta}(c+k, d+l)$$

DP Design

- We use dynamic programming (DP) to obtain an optimal adaptive treatment allocation sequence
- Optimal in the sense of maximising the expected total number of successes in the trial
- Specifically, we use backward recursion algorithm
- Let $\mathcal{F}_m(i, j, k, l)$ be the expected total number of successes under an optimal policy after m patients
- If $m = n$, there is nothing to do: $\mathcal{F}_n(i, j, k, l) = 0$ $\forall i, j, k, l$

Backward Recursion

- If $m = n - 1$ (one patient left):

  1. If treatment C, we compute the expectation

$$\mathcal{F}^C_{n-1}(i, j, k, l) = \frac{i}{i+j} \cdot 1 + \frac{j}{i+j} \cdot 0$$

  2. If treatment D, we compute the expectation

$$\mathcal{F}^D_{n-1}(i, j, k, l) = \frac{k}{k+l} \cdot 1 + \frac{l}{k+l} \cdot 0$$

- We wish to choose the optimal allocation such that

$$\mathcal{F}_{n-1}(i, j, k, l) = \max\{\mathcal{F}^C_{n-1}(i, j, k, l),\ \mathcal{F}^D_{n-1}(i, j, k, l)\}$$

Computational Tractability

Publication	T	$T^{\max}$	SW, HW, RAM
Steck (1964)	25	N/A	N/A, UNIVAC 1105, 54 kB
Yakowitz (1969)	5	N/A	Fortran, N/A, N/A
Berry (1978)	100	N/A	Basic (?), Atari (?), N/A
Ginebra and Clayton (1999)	150	180	N/A, N/A, N/A
Hardwick et al. (2006)	100	200	N/A, N/A, N/A
Ahuja and Birge (2016)	96	240	N/A, Mac 4GB
Williamson et al. (2017)	100	215	R, PC, 16GB
Villar (2018)	100	N/A	Matlab, PC, N/A
Kaufmann (2018)	70	N/A	N/A, N/A, N/A

Computational Tractability

Publication	T	$T^{\max}$	SW, HW, RAM
Steck (1964)	25	N/A	N/A, UNIVAC 1105, 54 kB
Yakowitz (1969)	5	N/A	Fortran, N/A, N/A
Berry (1978)	100	N/A	Basic (?), Atari (?), N/A
Ginebra and Clayton (1999)	150	180	N/A, N/A, N/A
Hardwick et al. (2006)	100	200	N/A, N/A, N/A
Ahuja and Birge (2016)	96	240	N/A, Mac 4GB
Williamson et al. (2017)	100	215	R, PC, 16GB
Villar (2018)	100	N/A	Matlab, PC, N/A
Kaufmann (2018)	70	N/A	N/A, N/A, N/A
This talk	4440	4440	Julia 1.0.1 & BB, PC, 32GB

Computational Tractability

Software	RAM	$T = 60$	$T = 120$	$T = 180$	$T = 240$	$T = 300$	$T^{\max}$
Julia 0.6.2 & ad hoc	12 GB	2sec	22sec	108sec	331sec	789sec	420
Julia 1.0.1 & ad hoc	12 GB	1sec	17sec	82sec	262sec	643sec	420
R & ad hoc	16 GB	1sec	12sec	59sec	191sec	N/A	240
Julia 1.0.1 & BB	31 GB	0.0036sec	0.046sec	0.23sec	0.73sec	1.6sec	1440
R & ad hoc	5 GB	1sec	6sec	26sec	84sec	209sec	420
Julia 1.0.1 & BB	31 GB	0.0040sec	0.056sec	0.27sec	0.91sec	2.8sec	4440

Runtime

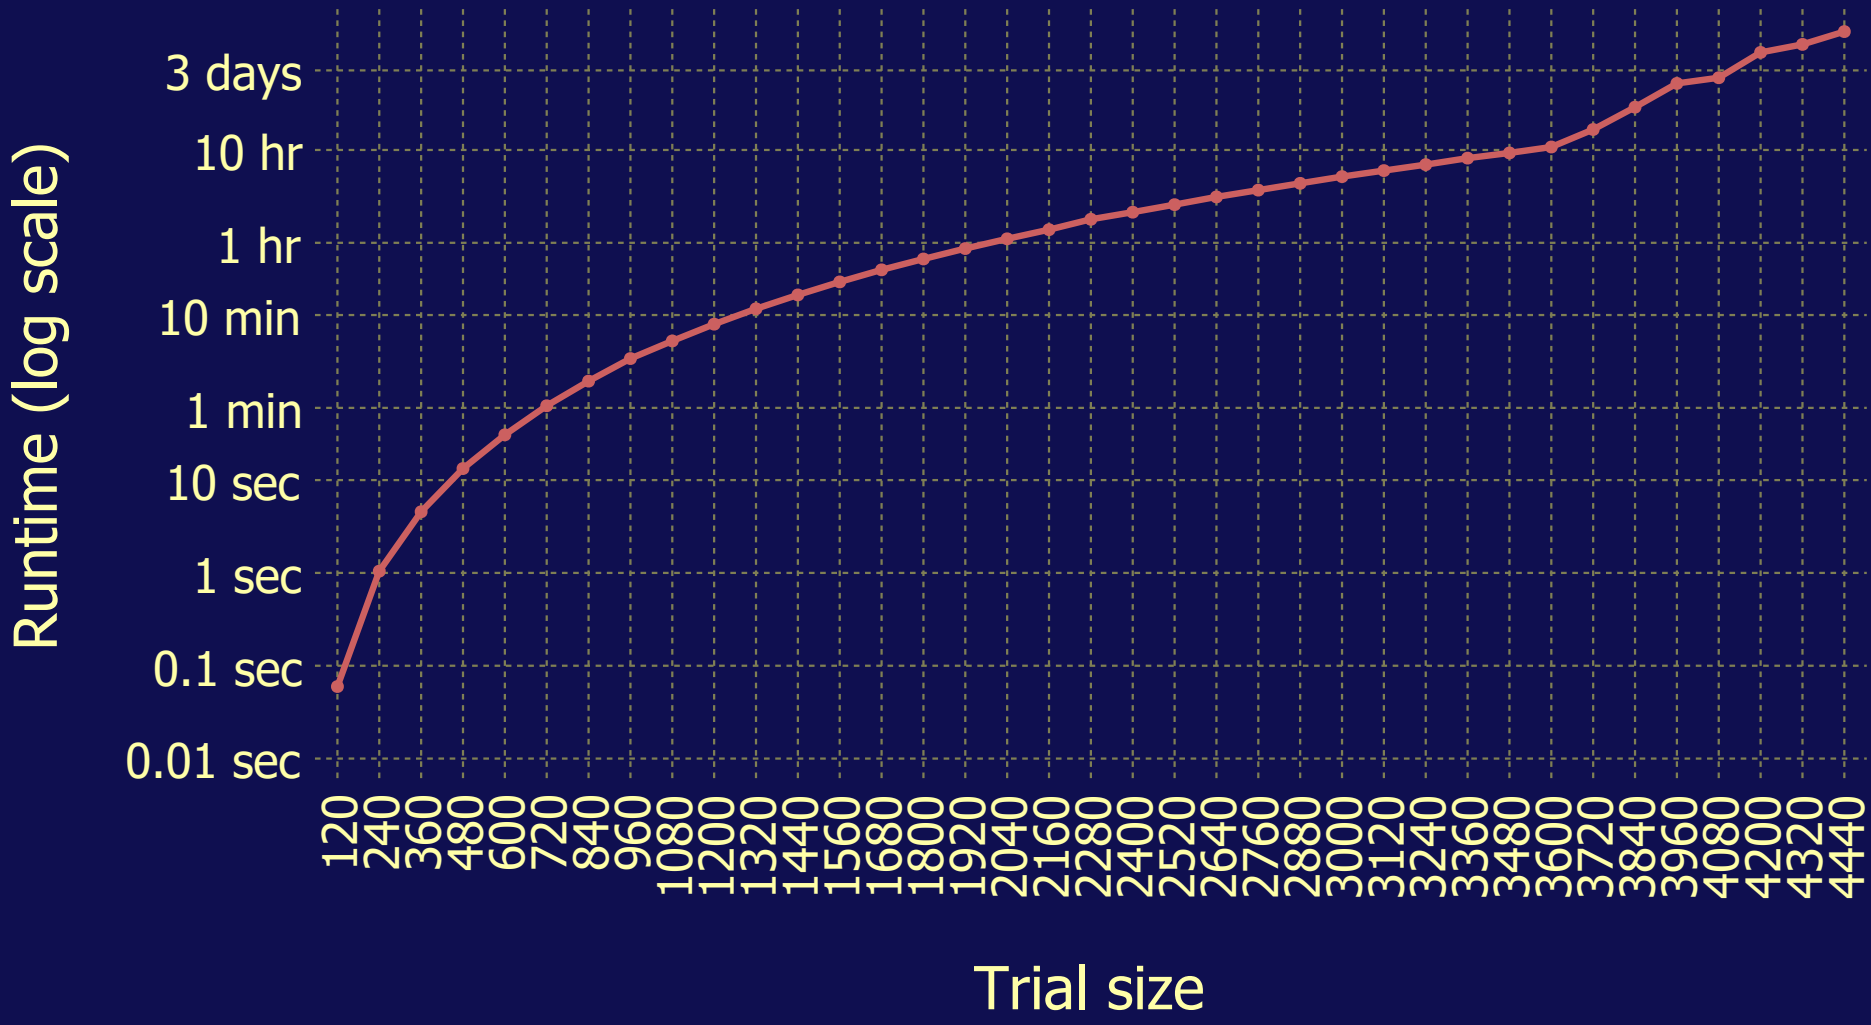
Runtime (log scale)
3 days
10 hr
1 hr
10 min
1 min
10 sec
1 sec
0.1 sec
0.01 sec
120
240
360
480
600
720
840
960
1080
1200
1320
1440
1560
1680
1800
1920
2040
2160
2280
2400
2520
2640
2760
2880
3000
3120
3240
3360
3480
3600
3720
3840
3960
4080
4200
4320
4440
Trial size

Memory Requirement

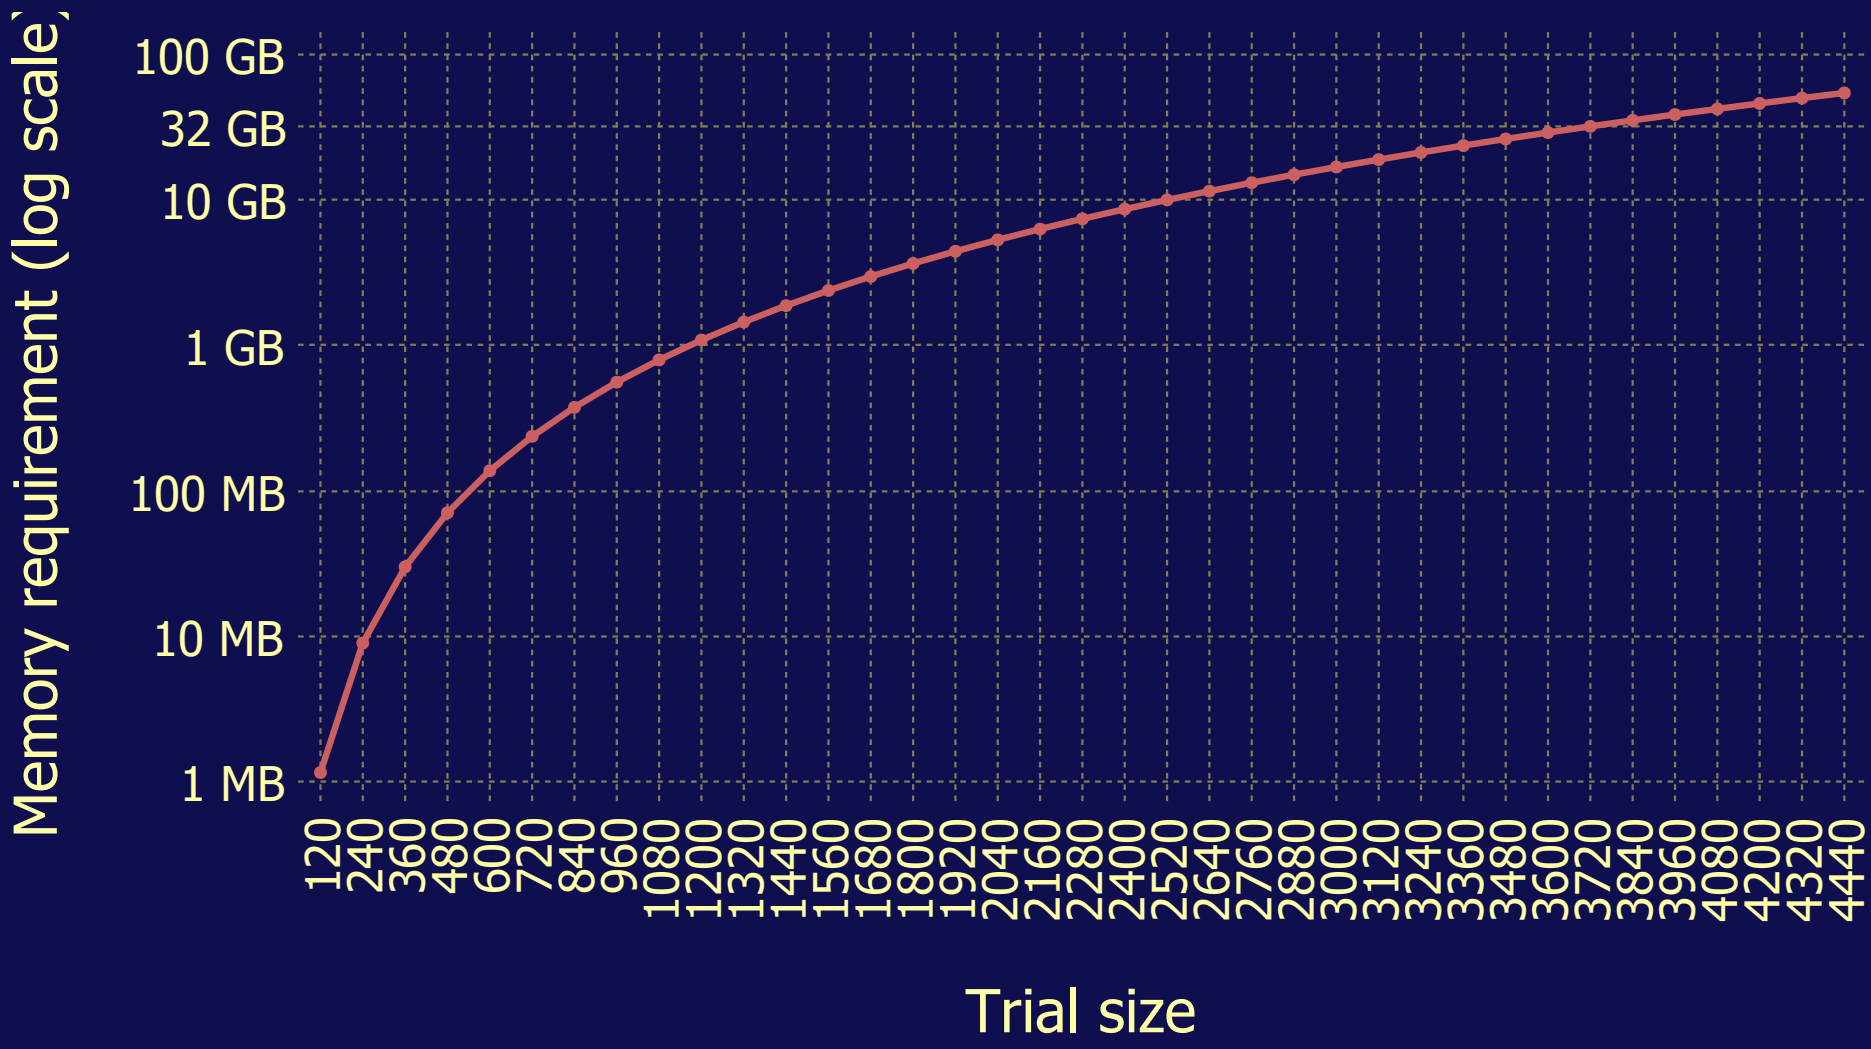

Bayes Proportion of Successes

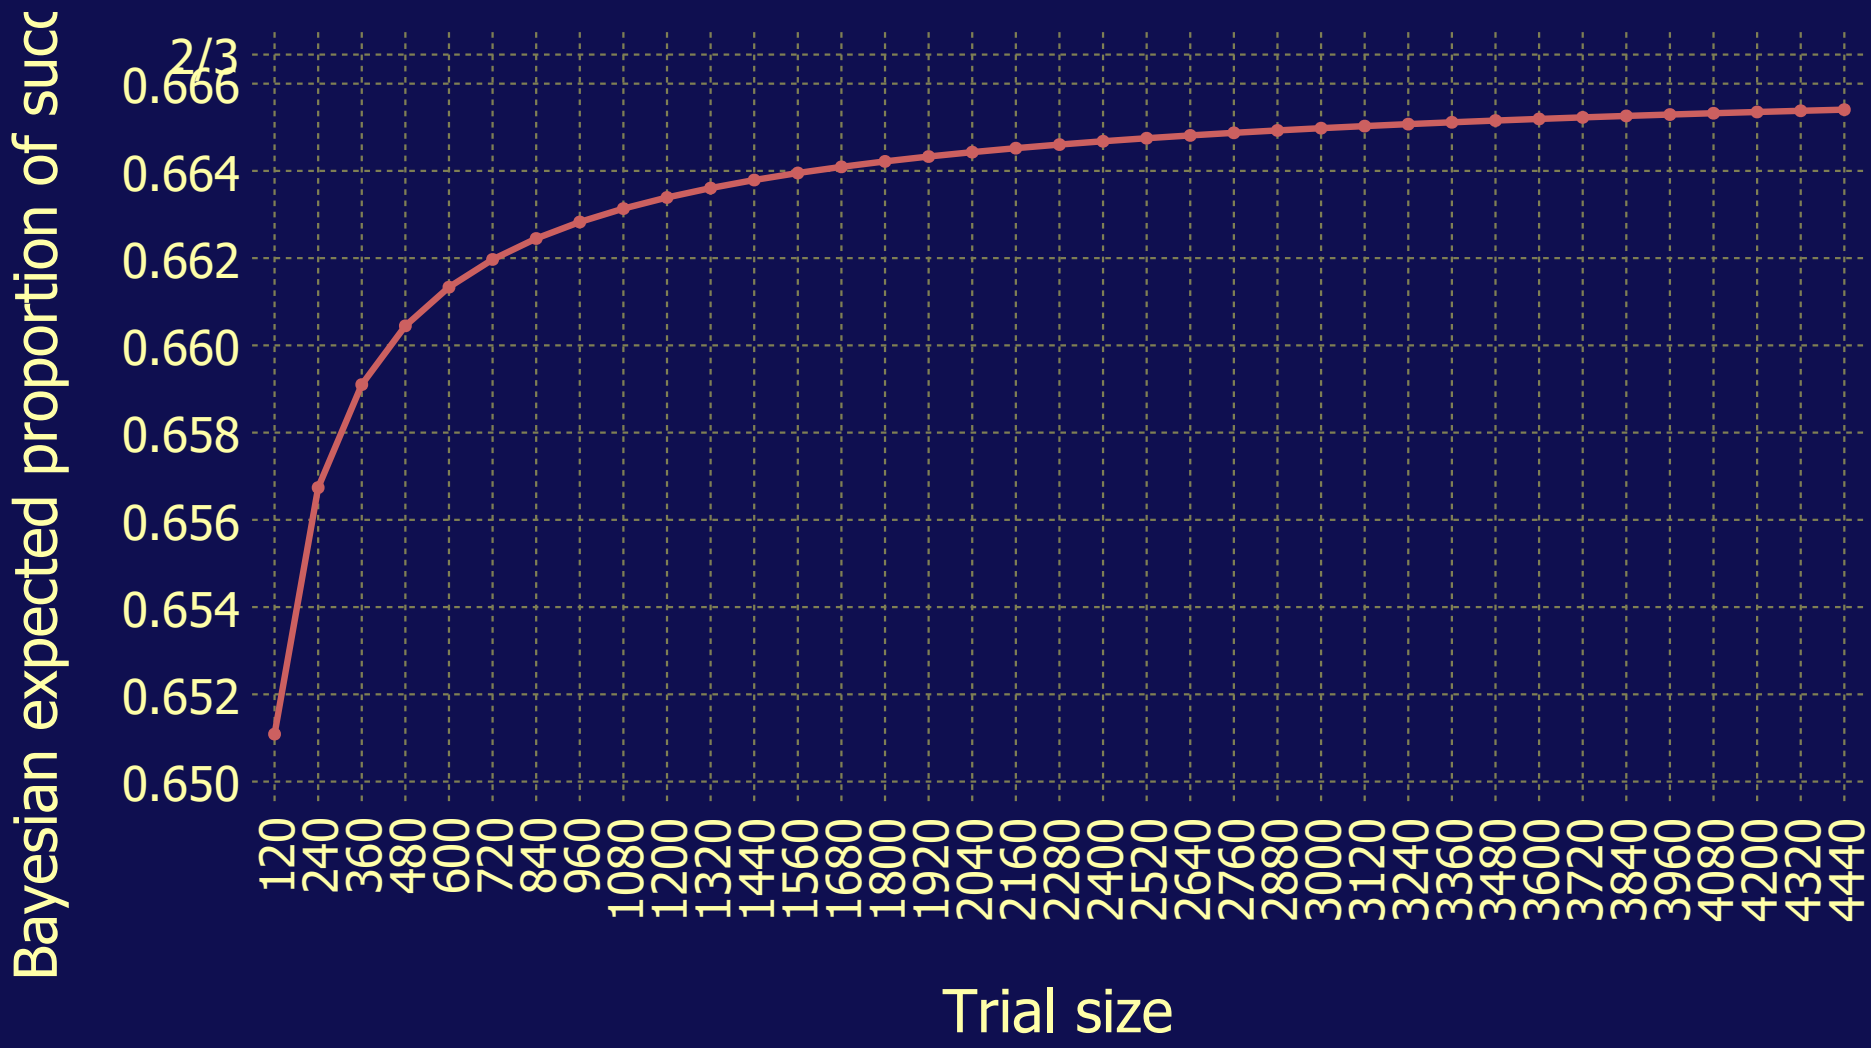

Bayes Regret Number of Successes

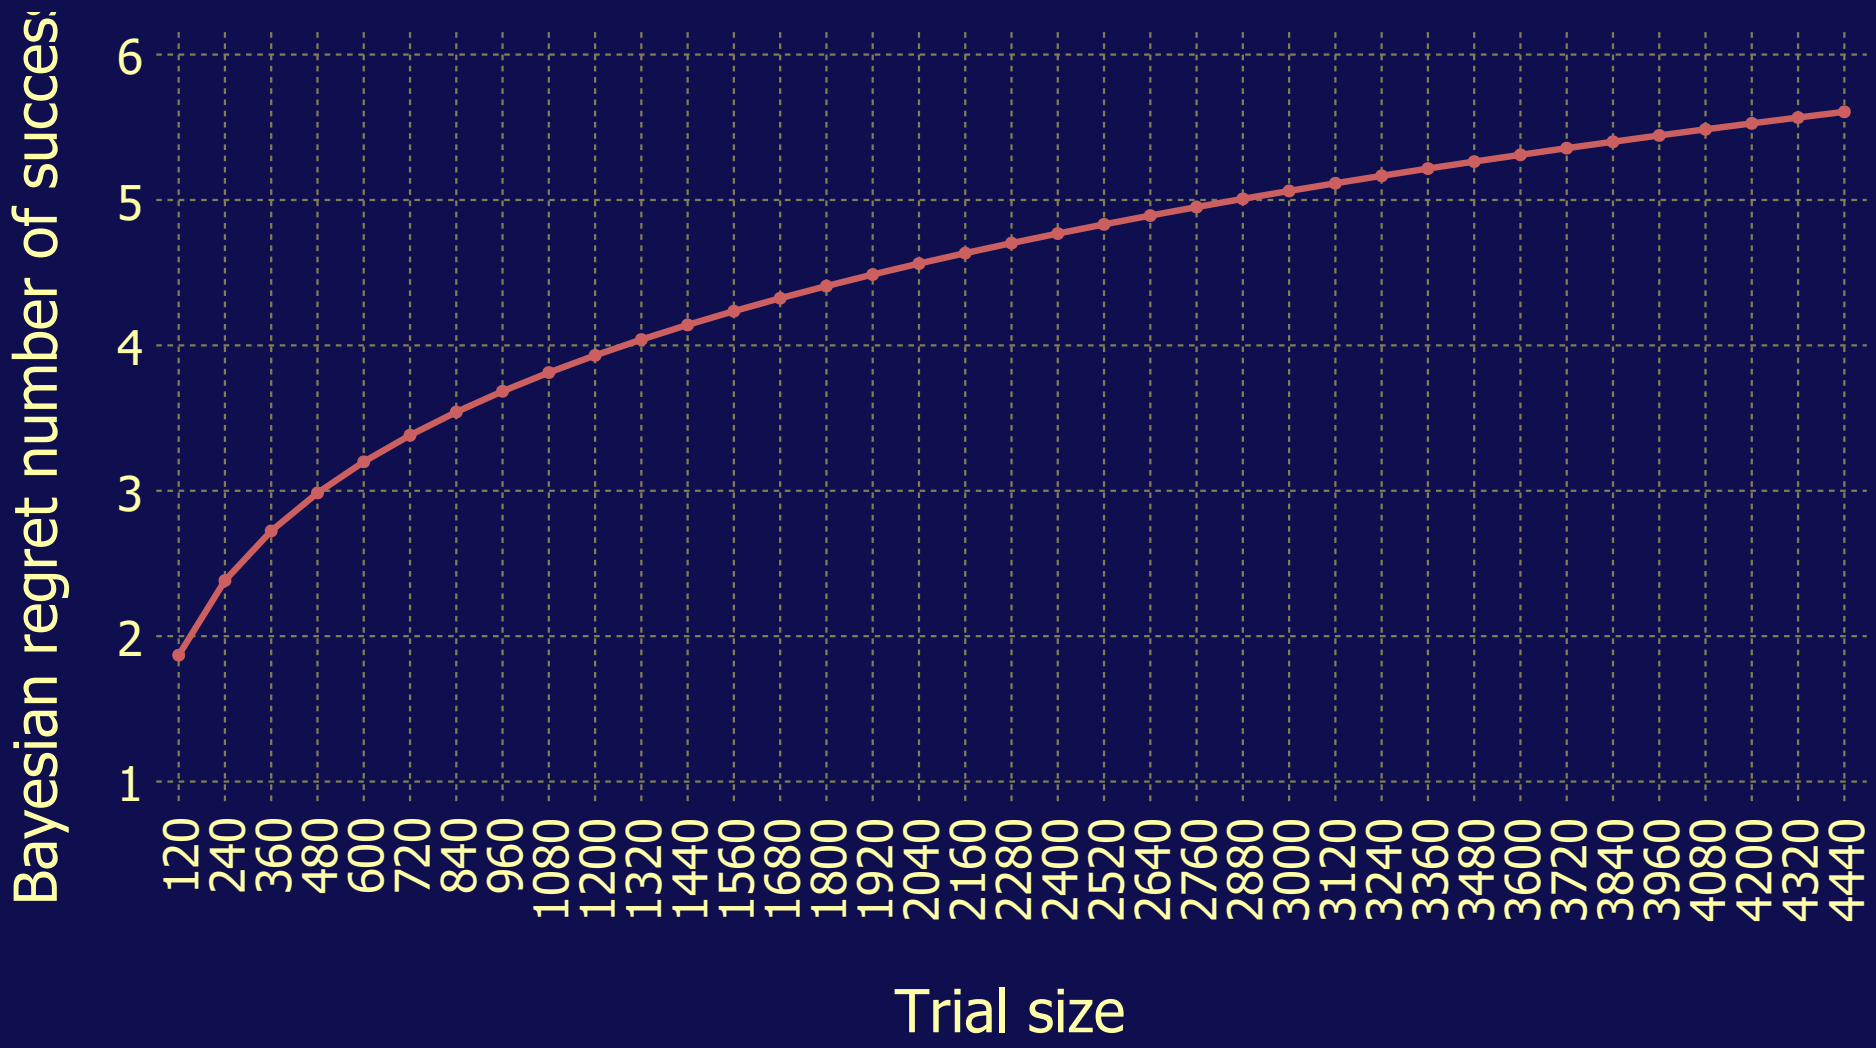

Other Designs

- For any other design, backward recursion can be used for evaluation instead of simulation
  - ▷ will be accurate (no simulation error)
- On the next slide we compare response-adaptive designs to 1:1, using frequentist regret
  - ▷ success probabilities $\theta_C = 0.7, \theta_D = 0.9$
  - ▷ BM/FM: Bayesian/frequentist myopic (naïve)
  - ▷ UCB: Upper Confidence Bound (machine learning)
  - ▷ BKG: Bayesian Knowledge Gradient (ranking & sel.)
  - ▷ BLFF: Bayesian Least Failures First

Frequentist Regret Number of Successes

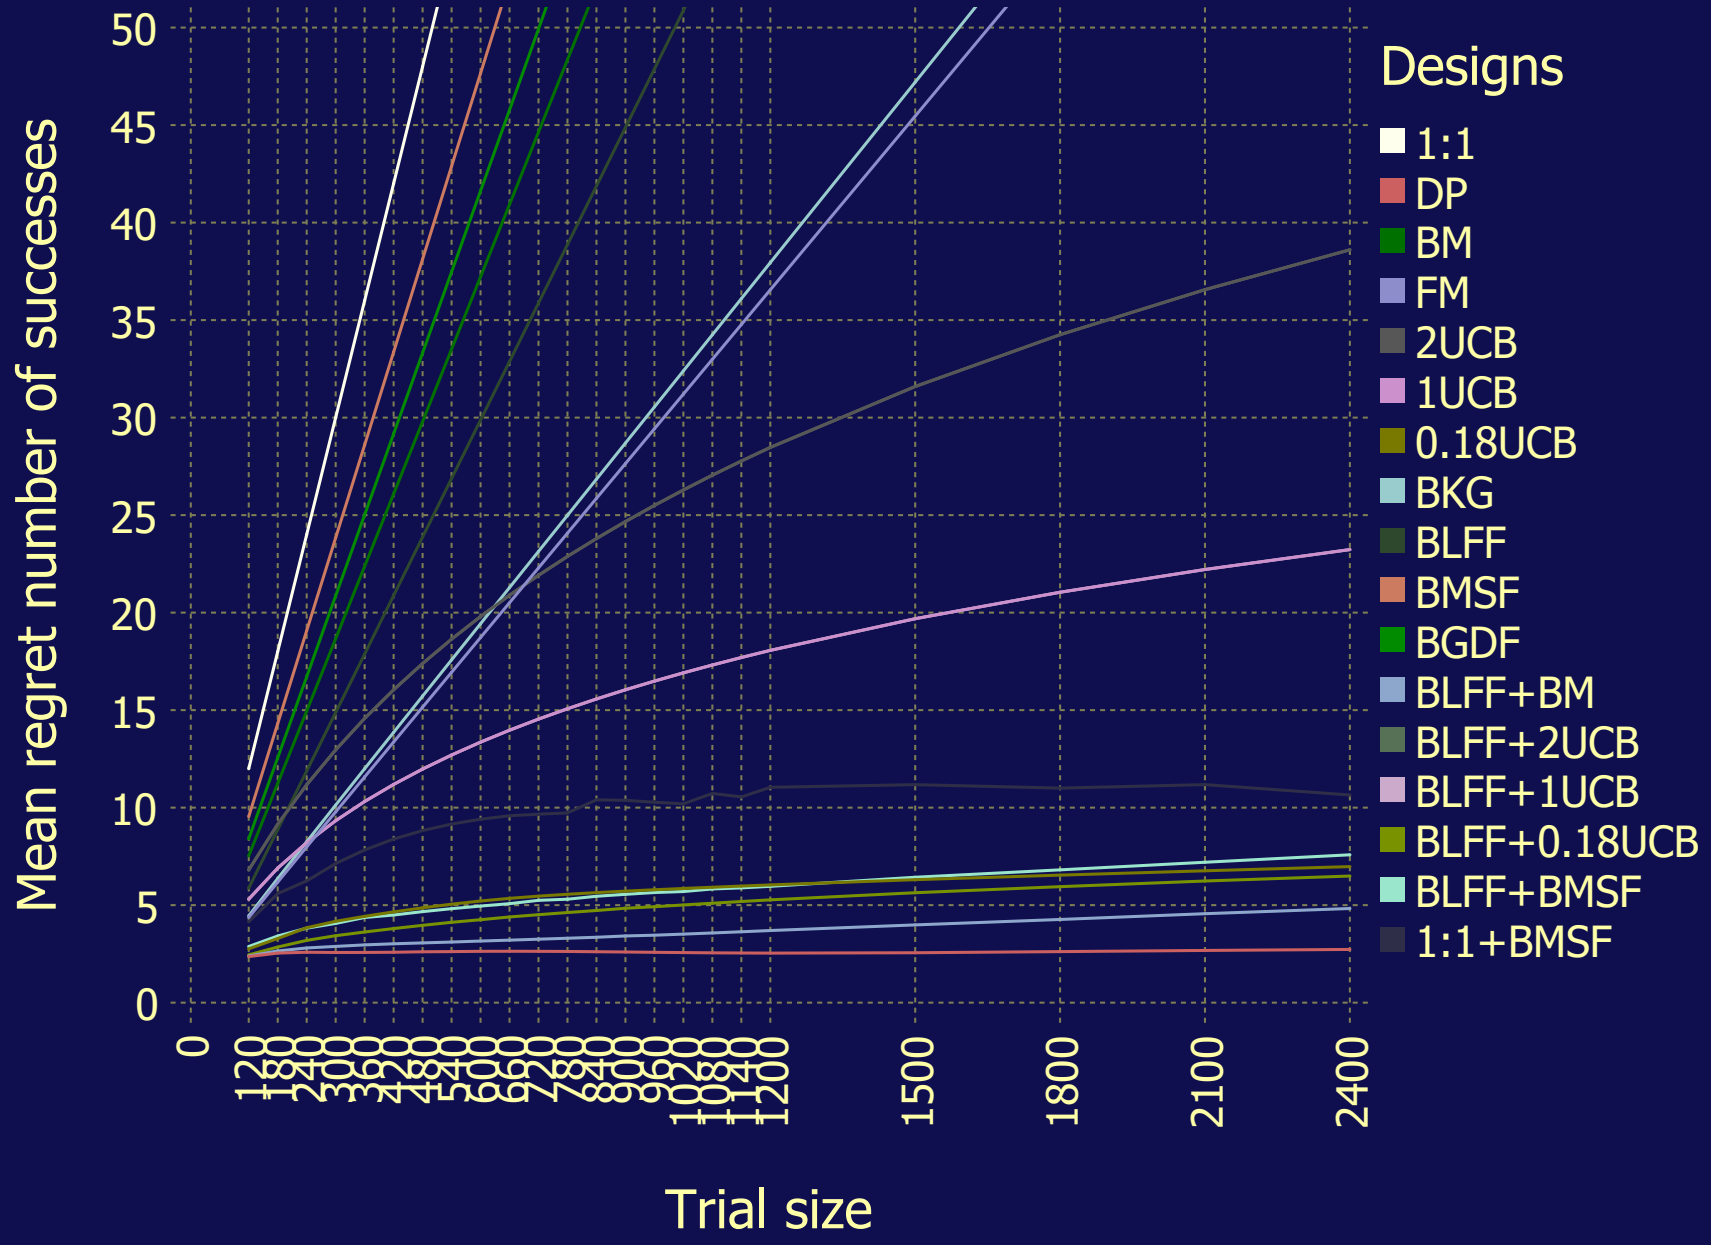

Thompson's Posterior Sampling Design

- Thompson 1933 proposed a heuristic:
  - ▷ randomise according to the posterior probability of being the best arm
- This can be done by exact calculation or by sampling
- Recently, several trials have been designed in this way
  - ▷ Don Berry (MD Anderson) and Berry Consultants
  - ▷ e.g. I SPY-2, GBM Agile
- Several recent papers by a group at Harvard
- "Easy" to use, but quite suboptimal

Conclusion about Designs

- Thompson's posterior sampling is myopic
  - ▷ randomisation decisions are bases on the assumption that the next patient is the last one
- Equal Randomization is utopic
  - ▷ randomisation decisions are bases on the assumption that there is an infinite number of future patients
- Optimal design is not myopic, not utopic
  - ▷ randomisation decisions take into account the remaining trial size and the after-trial population size
  - ▷ provides a significantly higher health benefit

Optimal Design

- The optimal design has not been implemented in practice yet
  - ▷ will you be the first?

Work in Progress

- I am working on a Julia package BinaryBandit (BB) to evaluate the operating characteristics of designs
  - ▷ more general settings will be considered
  - ▷ coding work for several years!
- With F. Williamson and T. Jaki, we are studying the (randomized) DP design if there are delayed responses
- With J. Wason, we are looking at when it is optimal to add in a novel treatment to a platform trial

Thank you for your attention

Group on Optimal Adaptive Learning (G.O.A.L.)
Lancaster University
http://www.lancaster.ac.uk/staff/jacko/goal/

References

Ahuja, V. and Birge, J. R. (2016). Response-adaptive designs for clinical trials: Simultaneous learning from multiple patients. *European Journal of Operational Research*, 248:619–633.

Berry, D. A. (1978). Modified two-armed bandit strategies for certain clinical trials. *Journal of the American Statistical Association*, 73:339–345.

Ginebra, J. and Clayton, M. K. (1999). Small-sample performance of Bernoulli two-armed bandit Bayesian strategies. *Journal of Statistical Planning and Inference*, 79(1):107–122.

Hardwick, J., Oehmke, R., and Stout, Q. F. (2006). New adaptive designs for delayed response models. *Journal of Statistical Planning and Inference*, 136:1940–1955.

Kaufmann, E. (2018). On Bayesian index policies for sequential resource allocation. *The Annals of Statistics*, 46(2):842–865.

Steck, R. (1964). A dynamic programming strategy for the two machine problem. *Mathematics of Computation*, 18(86):285–291.

Villar, S. S. (2018). Bandit strategies evaluated in the context of clinical trials in rare life-threatening diseases. *Probability in the Engineering and Informational Sciences*, 32:229–245.

Williamson, S. F., Jacko, P., Villar, S. S., and Jaki, T. (2017). A Bayesian adaptive design for clinical trials in rare diseases. *Computational Statistics and Data Analysis*, 113C:136–153.

Yakowitz, S. J. (1969). *Mathematics of Adaptive Control Processes*. New York, NY: North-Holland.