Innovative Stochastic Modelling and Optimisation for the Design of Modern Clinical Trials

Peter Jacko*

joint work with Sofia S. Villar, S. Faye Williamson, Thomas Jaki, James Wason

Queueing Theory Workshop
Obergurgl, December 18–20, 2017

*Dept. of Management Science, Lancaster University, UK

Traditional Clinical Trials

- Two treatments: control (existing one) and
  novel (not approved yet)

- Is the novel treatment better than the control?
  - ▷ if not enough evidence, it will not be approved!

- The gold standard design: randomised controlled trial
  - ▷ 50% vs 50% fixed equal randomisation of T patients
  - ▷ avoids all 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

Clinical Trials

- Problems with 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

Multi-Armed Bandit Problem

Multi-armed Bandit Model

- Maximise healing of patients in the trial
  - ▷ optimally solving learning/earning trade-off
  - ▷ learning takes place during the trial
- The multi-armed bandit motivated by clinical trials
  - ▷ Thompson (Biometrika 1933), Robbins (1952), etc.
- Bandit models are a type of response-adaptive design
- Appropriate model: finite horizon
  - ▷ can't be optimally decomposed!

Bernoulli Two-Armed Bandit Model

- Finite horizon: T sequentially arriving patients
- Two-armed: treatment A or B for each patient
- Binary endpoints: success (1) or failure (0)
- Let X_t and Y_t denote patient t's response from treatment A and B respectively (for $t = 1, \ldots, T$). Then,

$$X_t \sim \text{Bernoulli}(1,\ \theta_A) \;\; \text{and} \;\; Y_t \sim \text{Bernoulli}(1,\ \theta_B),$$

where θ_A and θ_B are the unknown success probabilities of treatments A and B respectively

Bayesian Approach

- Beliefs $\widehat{\theta}_A$ and $\widehat{\theta}_B$ to be updated over the trial

- Prior Distribution for $k \in \{A, B\}$: $\widehat{\theta}_k \sim \text{Beta}(s_k^0,\ f_k^0)$ where we take $s_k^0 = f_k^0 = 1$ (uninformative)

- Posterior Distribution: After observing s_k successes and f_k failures on treatment k, the posterior distribution is represented by another Beta distribution (by conjugacy)

$$\widehat{\theta}_A \sim \text{Beta}(s_A^0 + s_A, f_A^0 + f_A)$$

Optimal Design using DP

- 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 induction algorithm
- Let $\mathcal{F}_m(s_A, f_A, s_B, f_B)$ be the expected total number of successes under an optimal policy after m patients
- Using 4-dimensional state space (T^4)

Optimal Design using CRDP

- Practical problem? deterministic, underpowered, etc.

Example. $T = 75, \ell = 0.15T, \theta_A = 0.2, \theta_B = 0.6$

p	Bias	MSE	Power	EPS	On sup
60%	0.002	0.005	0.935	43.7%	59.1%
70%	0.002	0.007	0.910	47.3%	68.2%
80%	0.005	0.009	0.834	50.9%	77.2%
90%	0.008	0.013	0.724	53.6%	84.0%

Simulation Results: Designs Comparison

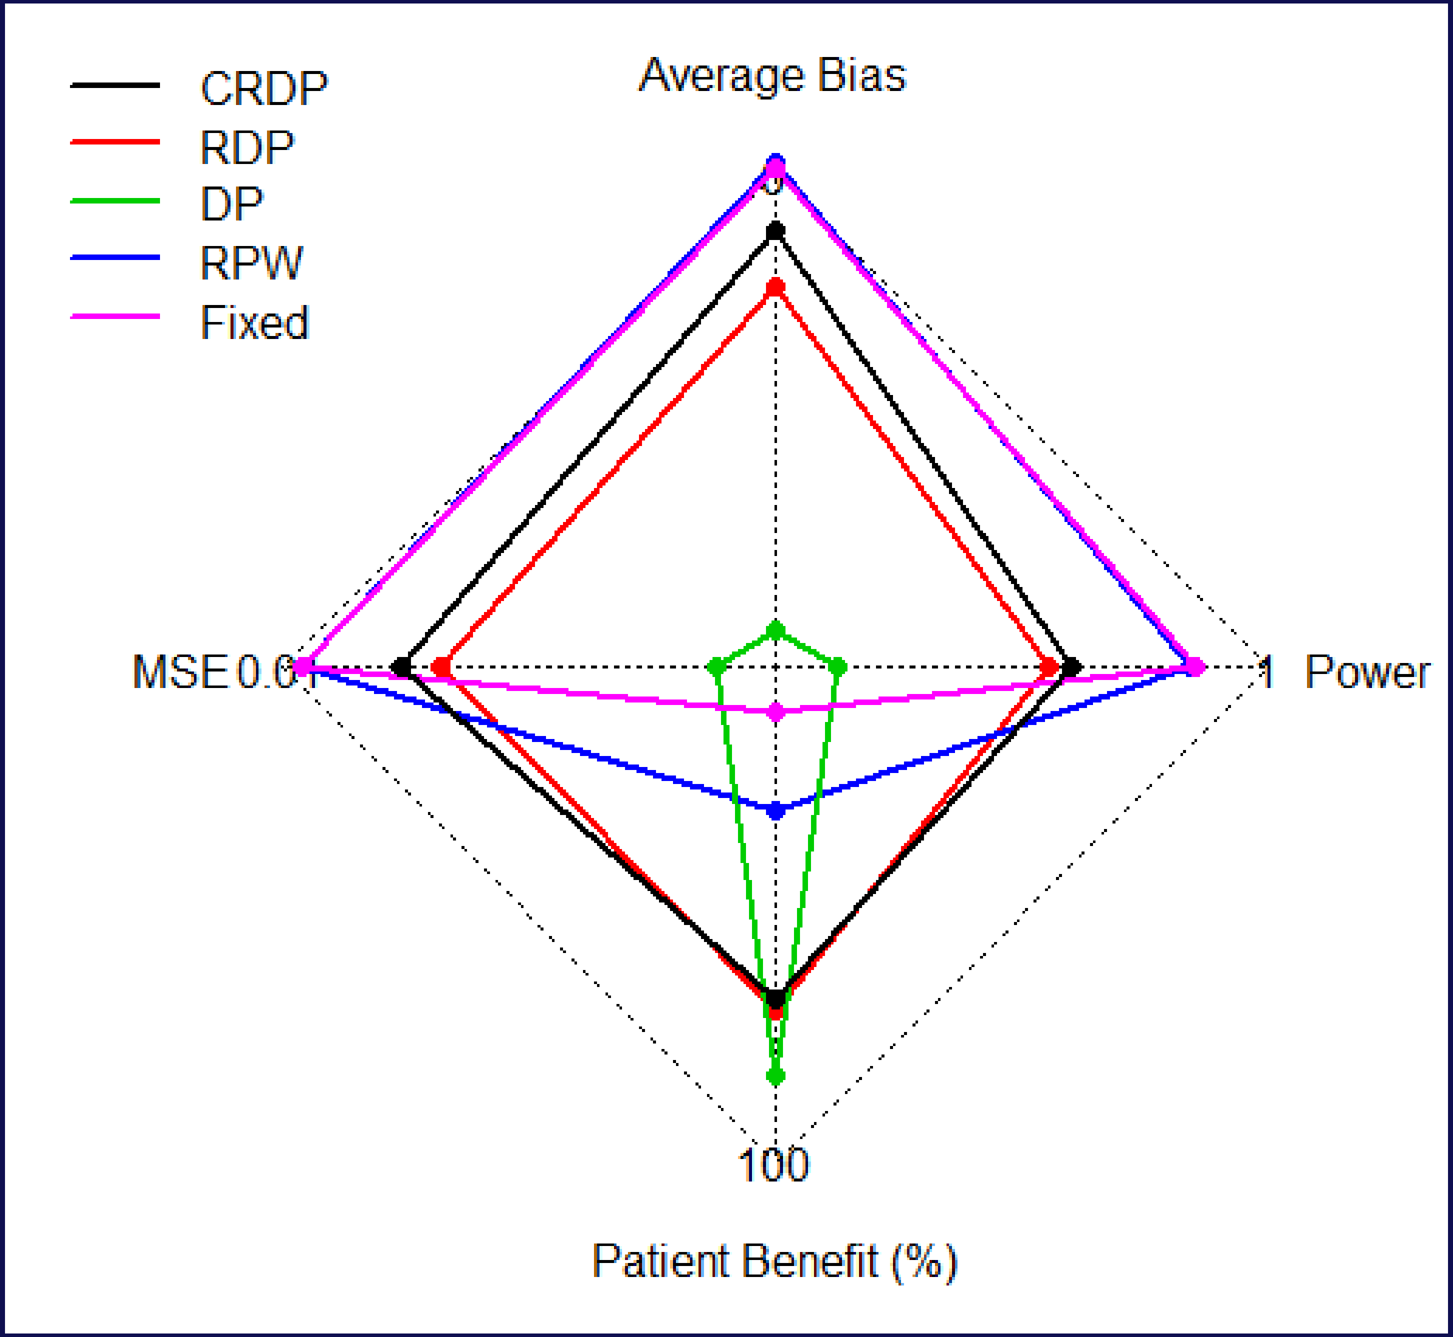

Conclusion

- We need to talk to statisticians and clinicians about bandit models
  - ▷ give me randomisation probability and desired power
  - ▷ I tell how to randomise to heal most patients
- Trials of the 21st century
  - ▷ stratification of patients to achieve personalised treatments
  - ▷ involvement of patient opinions in drug development
  - ▷ decision-making based on small samples

Julia Programming Language

- My experience for MDPs and DP: huge improvements
  - e.g. R: could run up to $T = 200$: 25min, 12GB array
  - Julia: could run up to $T = 360$: 1min, 12GB array
  - Julia: $T = 200$: 10sec, 3GB array

- More improvement possible with a few tricks
  - Julia: could run up to $T = 900$: 20min, 25GB array
  - on a laptop with 16GB RAM!!!

Thank you for your attention

...and see you in Lancaster...

- The 7th meeting of the EURO WG on StochMod
- 13–15 June 2018, Lancaster Management School
- Become member at www.stochmod.eu — it's free!

...or in Rotterdam...

- The (2nd) Workshop on Multi-Armed Bandits and Learning Algorithms
- 24–25 May 2018, Rotterdam School of Management