Designing Bayesian clinical trials to enhance vaccine policy recommendations

Michael Dymock et al.

Acknowledgements

- Supervisors:
  - Kevin Murray (UWA)
  - Julie Marsh (The Kids/AHI)
  - Tom Snelling (USyd/AHI)
- Advisory Panel:
  - Kate Lee (MCRI)
  - Rob Mahar (UniMelb)
- More at *michaeldymock25@github.io*

THE UNIVERSITY OF WESTERN AUSTRALIA

The KIDS RESEARCH INSTITUTE AUSTRALIA

ADAPTIVE HEALTH INTELLIGENCE

EVIDENCE IN ACTION

Outline

- Clinical research from first principles
  - Why do we do it? How do we do it?
- The value-driven adaptive design
- RSV case study

Why do clinical research?

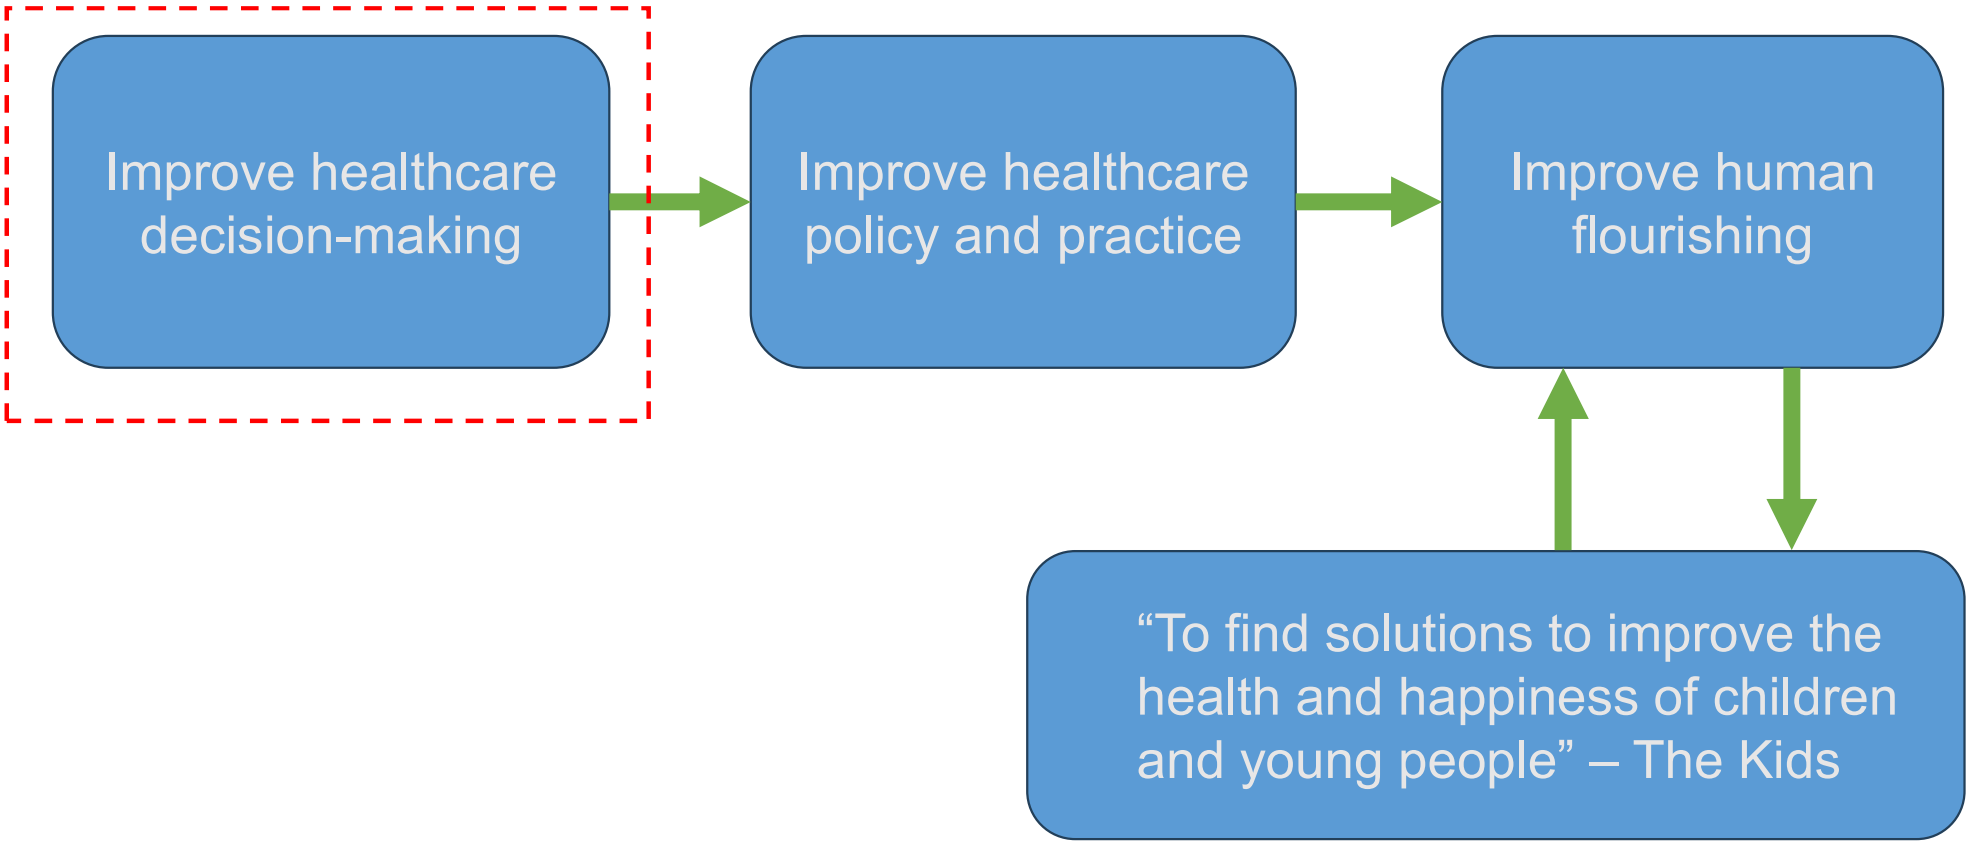

Is there a problem?

- Results from clinical research only slowly and inconsistently inform policy and practice
- Many reasons for this (of which some are statistical)
- Clinical research is typically designed for drug/vaccine registration
- But it *should** **prioritise** informing the **decision-making** of consumers, clinicians and policy-makers (i.e., improve healthcare decision-making)

Hypothesis testing

- A scientific hypothesis is distilled into a statistical hypothesis test
  - Suppose we want to estimate the treatment effect β
  - E.g., **A**: treatment has **no** effect vs **B**: treatment has a **positive** effect
- Scientists (humans) want to make a declaration: **A** is true, or **B** is true
  - This is a **truth claim** - it may be wrong
  - Frequentists and Bayesians do this

		Does the treatment have a positive effect?	
		Yes	No
Did I declare that the treatment has a positive effect?	Yes	Power	Type One Error
	No	Type Two Error	

Traditional designs

Everything focuses on this comparison:

Our decision rules (i.e., stop or continue recruitment)

Our sample size (at least in theory)

Our determination of trial “success”

Our decision on how to publish

Does this suit clinical research?

- We aim to improve healthcare **decision-making**
- Usually, we specify a single effect measure for a single primary outcome and make a single declaration
- But! A decision-maker will consider multiple outcomes (e.g., efficacy and safety) and cares about the effect sizes and their uncertainty!

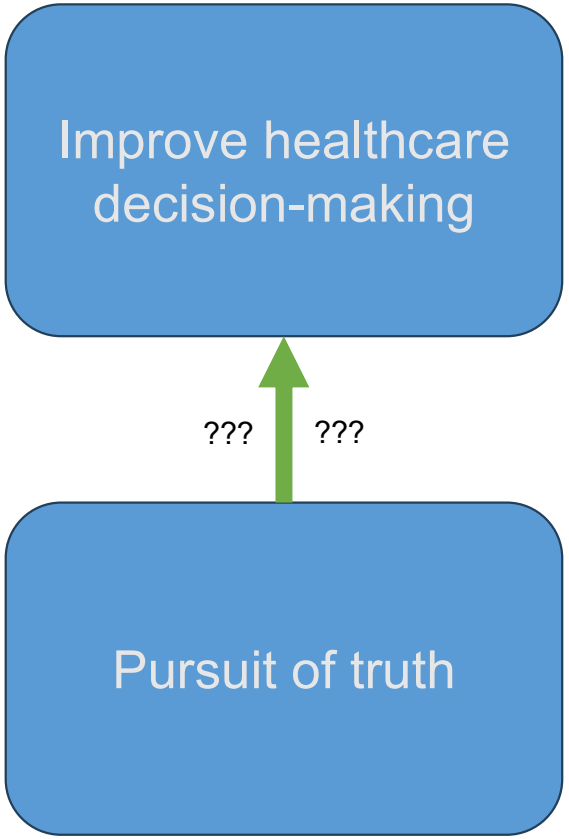

The role of science in the pursuit of truth

- "Science" is a tool we can use to uncover truths about the universe
- It is good to do science to better understand phenomena
- It will lead to improved healthcare decision-making and human flourishing
- But! We have finite resources!
- We should pursue the **right** amount of truth to the **right** questions

What if we did it differently?

- What if we **knew** what the decision-maker needed and designed a trial to answer this question **directly**?
- Suppose we asked decision-makers to specify a function that represents their decision-making process
- What if we designed a trial to collect **just** enough information to **sufficiently** inform the decision?
- No longer concerned with type one error and power because we have no interest in making declarations about the value of the effect parameter
- Could we **"bridge"** the gap (abyss) between clinical research and translation?

Specifying the decision-making process

Choose the decision option that that maximises the net benefit function:

$$\text{argmax}_d\big(\text{NB}(d, \boldsymbol{\Theta})\big)$$

E.g., if $\text{INB}(\boldsymbol{\Theta}) = \text{NB}(A, \boldsymbol{\Theta}) - \text{NB}(B, \boldsymbol{\Theta}) > 0$
we choose decision A instead of decision B – easy!

But what about the uncertainty?

- Our current understanding of $\boldsymbol{\Theta}$ is a distribution, not an exact value
- For example:

$$\boldsymbol{\Theta} = \{\theta_A, \theta_B\}$$

$$\theta_A \sim N(1, \sigma)$$

$$\theta_B \sim N(0, \sigma)$$

$$\mathrm{INB}(\boldsymbol{\Theta}) = \theta_A - \theta_B$$

- We can visualise the distribution of the INB
- If **forced** to decide now we could use the expected value $E_{\boldsymbol{\Theta}}[\mathrm{INB}(\boldsymbol{\Theta})]$

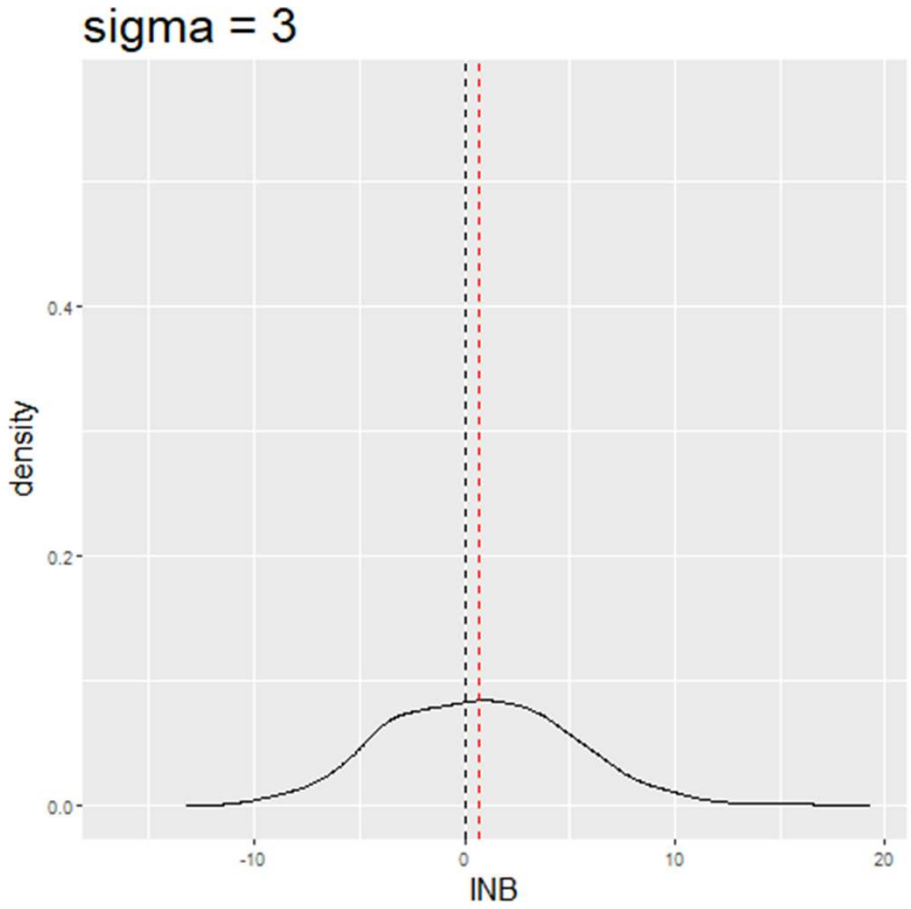

What is the value of reducing uncertainty?

- Is it worthwhile reducing the uncertainty by collecting information?
- We propose to collect data $\boldsymbol{X}$ to inform $\boldsymbol{\Theta}$ (e.g., by conducting a clinical trial)

$$\text{EVSI} = E_{\boldsymbol{X}}[E_{\boldsymbol{\Theta}|\boldsymbol{X}}[\max_{d} \text{NB}(d, \boldsymbol{\Theta})]] - \max_{d} E_{\boldsymbol{\Theta}}[\text{NB}(d, \boldsymbol{\Theta})]$$

Expected value of decision **after** collecting data

Expected value of **current** decision

- Define $\text{ENBS} = \text{EVSI} - \eta$, where η is some prespecified cost, and proceed if $\text{ENBS} > 0$

What does this look like?

But the EVSI depends on $\boldsymbol{\Theta}$ and $\boldsymbol{X}$!

These are unknown and require assumptions

Should we revise the EVSI as the trial proceeds?

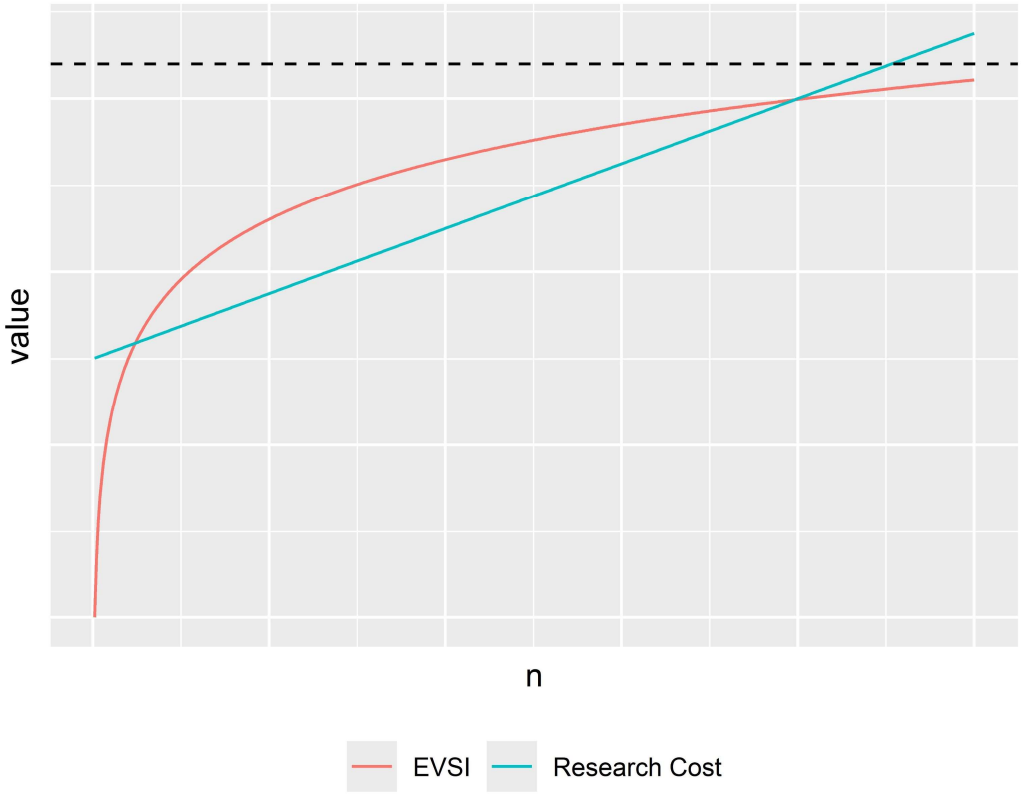

Notation	Definition
$\boldsymbol{p} = \{p_1, p_2, \ldots, p_D\}$	Current implementation (proportion) of decision options

Expected net benefit of sampling

$$\mathrm{ENBS}^{j} = E_{\boldsymbol{\Theta}|\boldsymbol{x}_{1:\bar{n}_j}}\left[\sum_{d=1}^{D} p_d \times \mathrm{NB}\left(d, \boldsymbol{\Theta}, t_j, t_{j+1}\right)\right] +$$

The value-driven adaptive design

- What if we collected data until the **cost of further recruitment** outweighed the **incremental benefit to our decision-making**?
- Recalculate the ENBS at interim analyses to make stop-go decisions
- Stop once $\mathrm{ENBS} < 0$
- Choose the best decision option and report out the results

Some notes

- Under this setup the ENBS is underestimated because we are only looking one step ahead (i.e., the calculation assumes we collect data and STOP at the next analysis)

- It is not computationally feasible to remedy the underestimation but an upper bound can be estimated using the EVPI

- There are actually three choices one could make at each analysis:

  A) Recruit a further n_{j+1} participants and proceed to the next analysis

  B) Stop recruitment and wait for participants to complete follow-up before choosing

  C) Stop recruitment and choose now

Choosing interim sample sizes

- The design is flexible to any policy for choosing interim sample sizes
- Usually based on practical constraints (e.g., anticipated recruitment, statistical resources, outcome ascertainment)
- But one may choose sample sizes strategically

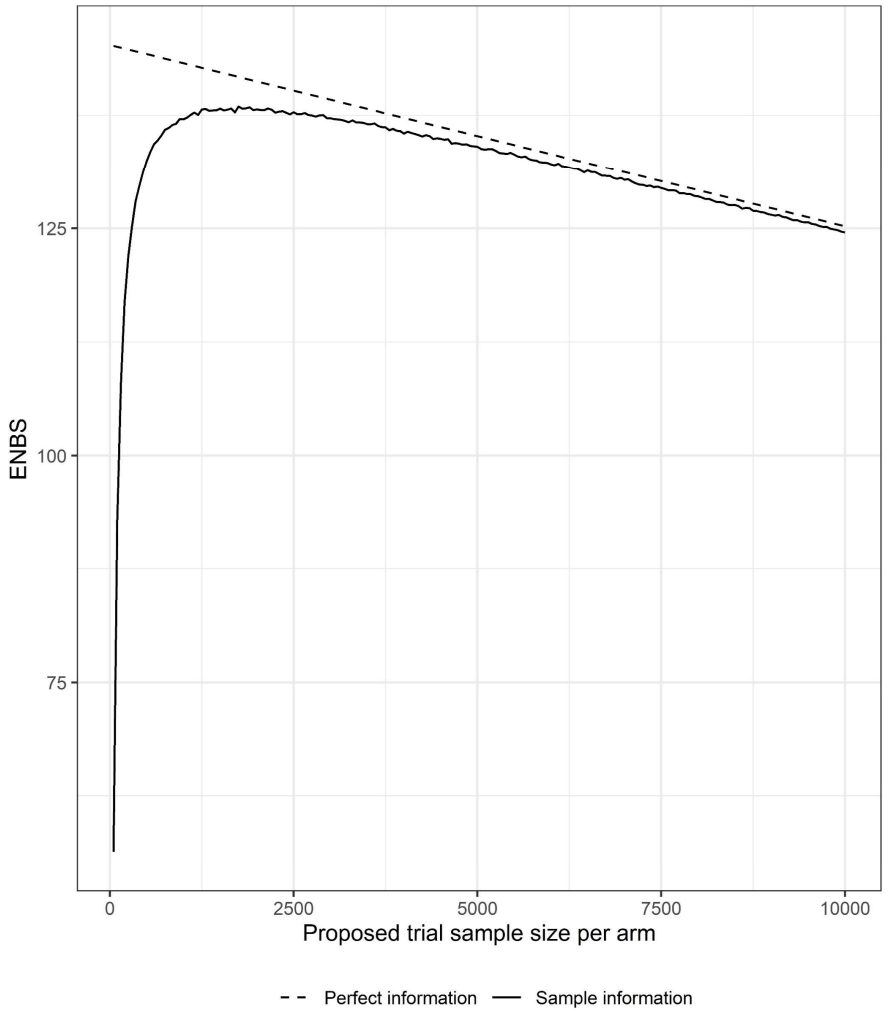

Case Study

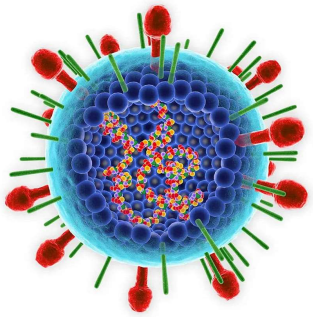

- Respiratory syncytial virus (RSV) infection accounts for approximately 3.6 million hospitalisations and over 100,000 deaths each year, globally
- In Australia it is unknown whether infant immunoprophylaxis (II) $(d = 1)$ or maternal vaccination (MV) $(d = 2)$ will be more cost-effective
- Assume II is current implemented practice with 100% uptake (i.e., $\boldsymbol{p} = \{1,0\}$)
- Interested in the trade-off between the cost of the strategies and the effectiveness in preventing medically attended RSV events (MA-RSV) in the first 12 months of life

The (simplified) health economic model

MV is more cost-effective than II if:

$$\mathrm{INMB}(\boldsymbol{\Theta}, \mathrm{t}_1, \mathrm{t}_2) > 0 \leftrightarrow p_{\mathrm{MV}} - p_{\mathrm{II}} < \frac{260}{\lambda}$$

We set $\lambda = \$5{,}200$ so that MV is more cost-effective if:

$$p_{\mathrm{MV}} - p_{\mathrm{II}} < 5\%$$

A clinical trial

- Randomise up to 1,000 mother-infant dyads (i_k) to each strategy $k \in \{\mathrm{II}, \mathrm{MV}\}$
- Dyads have 12-month MA-RSV outcome $x_{i_k} \in \{0,1\}$
- Analyses after every 250 dyads per strategy $(\mathrm{i.e.,}\ \ n_{1k} = n_{2k} = n_{3k} = n_{4k} = 250)$
  - One year to recruit each batch of participants and one year to ascertain outcomes
- Fixed start-up cost $\delta = \$1\ \mathrm{million}$ and per-dyad recruitment cost $\gamma = \$2{,}000$
- Weakly informative prior distributions: $p_k \sim \mathrm{Beta}(4,20)$
- Binomial likelihood: $\sum_{i_k=1}^{\bar{n}_{jk}} X_{i_k} \sim \mathrm{Binomial}\left(\bar{n}_{jk}, p_k\right)$

Two illustrative scenarios

1) The incremental effectiveness of II compared to MV is large
   - $p_{\mathrm{II}} = 0.10$ and $p_{\mathrm{MV}} = 0.18$
2) The incremental effectiveness of II compared to MV is small
   - $p_{\mathrm{II}} = 0.10$ and $p_{\mathrm{MV}} = 0.12$

- For both scenarios we estimate $\mathrm{ENBS}^0 \approx \$120\ \mathrm{million}$ using Strong et al.'s non-parametric regression method so the trials proceed

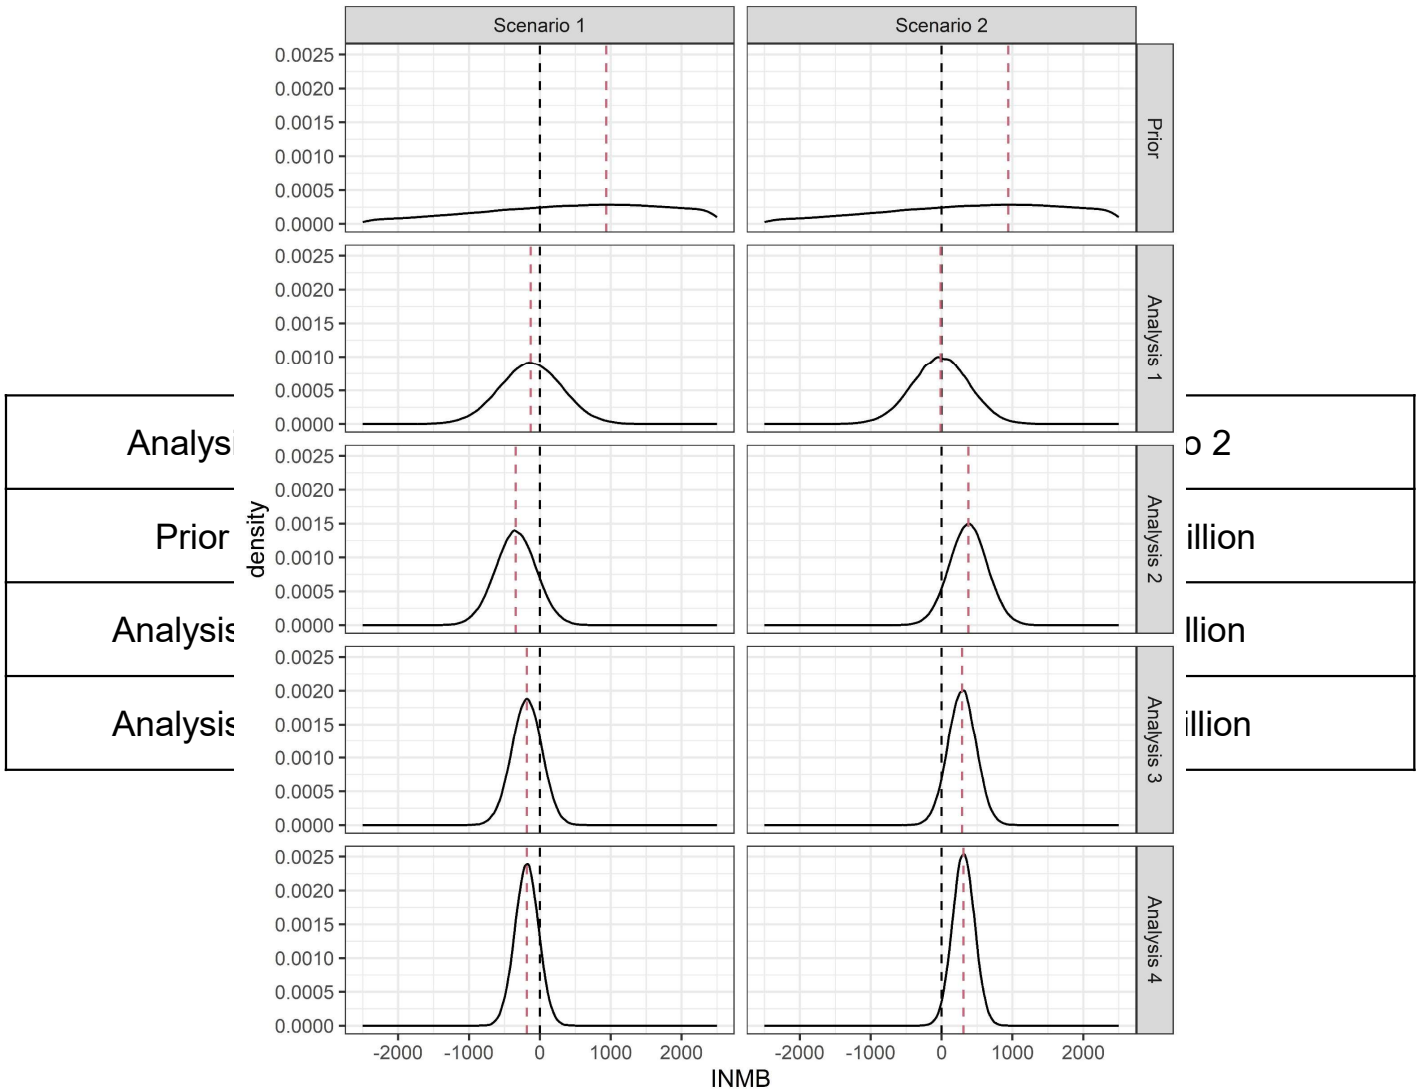
Scenario 1
Scenario 2
Prior
Analysis 1
Analysis 2
Analysis 3
Analysis 4
density
INMB
0.0025
0.0020
0.0015
0.0010
0.0005
0.0000
-2000
-1000
0
1000
2000

Current state and future directions

- Decision-theoretic methods for sequential designs have been developed that focus on estimating the mean INB directly
- More recently, Chick, Forster et al. have developed the value-based sequential design for a two-arm trial with a delay in outcome ascertainment
- Implemented by Flight, Brennan et al. retrospectively to UK funded trials
- R code implementation at https://github.com/michaeldymock25/htadelayR
- R code implementation of the value-driven adaptive design at https://github.com/michaeldymock25/ValueAdapt
- One day perhaps we could prospectively design a trial using these methods

Summary

Clinical research should* be designed to inform decision-making

The value-driven adaptive design is fundamentally different to traditional designs (adaptive or not) in its philosophy

- Not based on a hypothesis test (i.e., no type one error or power)

- Focused on the value of reducing a decision-maker's uncertainty

There are further complexities to consider (an adaptive design may not be appropriate, the health economic model can be extended, etc.)