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*

Outline

- Clinical research from first principles
  - Why do we do it? How do we do it?
- Vaccine policy recommendations in Australia
  - What is the process? What is the quality of evidence?
- A potential solution
  - The value-driven adaptive design

Clinical research from first principles

Why do we do it?

How do we do it?

Why do clinical research?

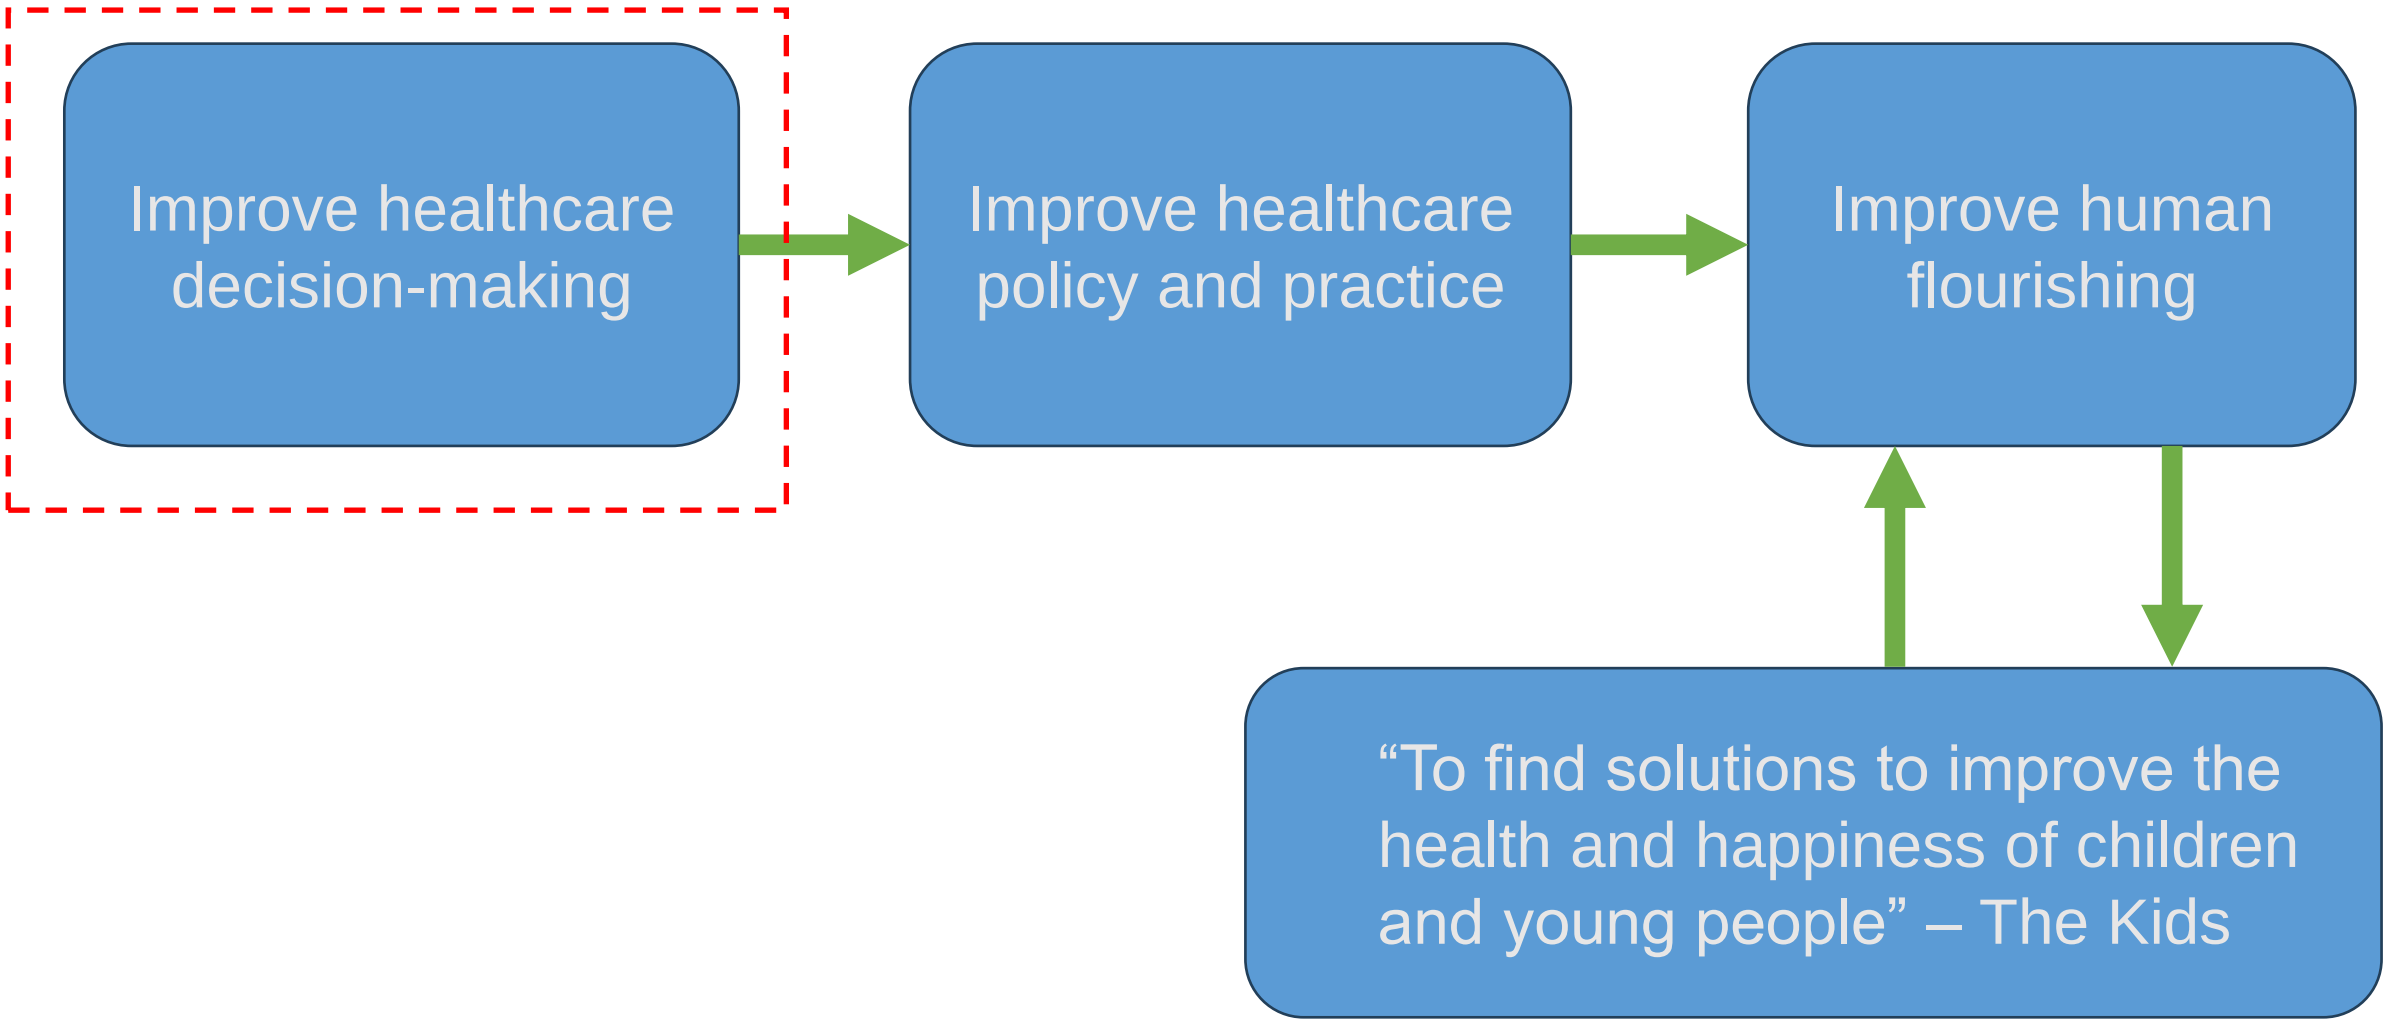

Is there a problem?

- 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!

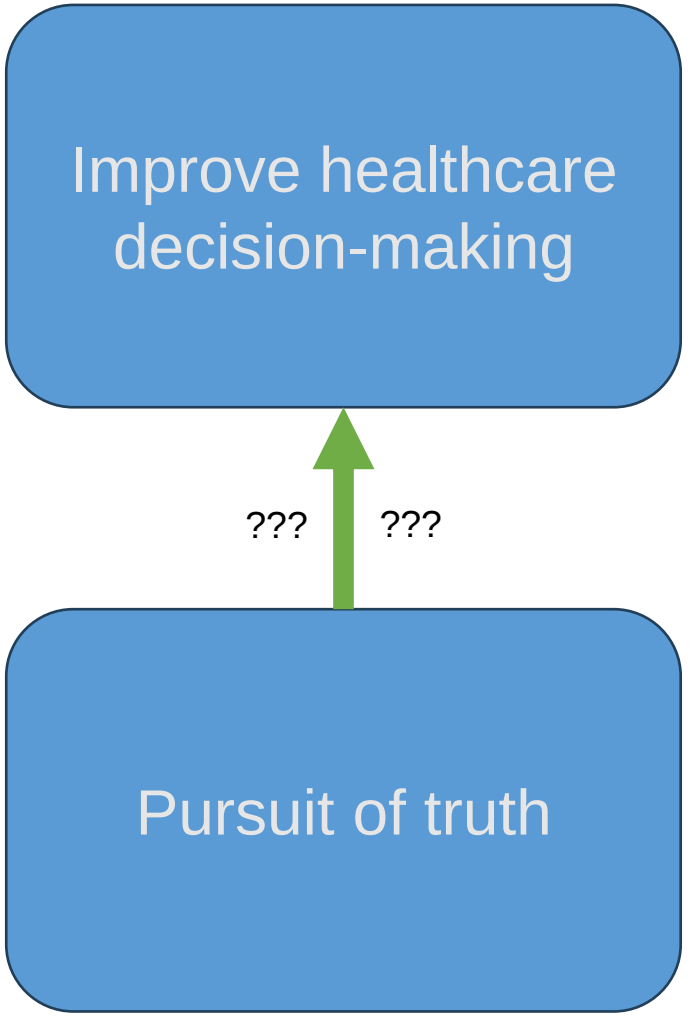

The role of science in the pursuit of truth

- “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

Vaccine policy recommendations in Australia

What is the process?

What is the quality of evidence?

Vaccine policy in Australia

- Applications are made to the Therapeutic Goods Administration (TGA)
- The Australian Technical Advisory Group on Immunisation (ATAGI) make recommendations to the Department of Health and Aged Care
- Vaccine policy recommendations approved by the National Health and Medical Research Council (NHMRC) are implemented in the Australian Immunisation Handbook

Australian Government
Department of Health and Aged Care

Australian Immunisation Handbook

Vaccine policy recommendations in Australia

- ATAGI's decision-making process is transparently documented
- GRADE method implemented by National Center for Immunisation Research and Surveillance (NCIRS) to develop recommendations
- Set the policy question (PICO)
- Identify literature (mostly clinical trials)
- Determine outcome-specific and overall GRADE
- Make recommendation (or no recommendation)

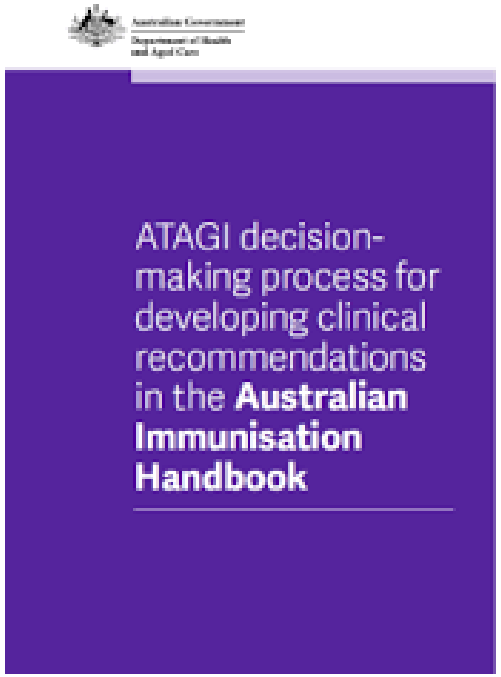

Grading of Recommendations Assessment, Development and Evaluation (GRADE)

- Used internationally by the NHMRC, World Health Organization (WHO) and the Advisory Committee on Immunization Practices (ACIP)
- Each outcome assessed for *risk of bias*, *inconsistency*, *indirectness* and *imprecision*
- Evidence **quality** classified as *very low*, *low*, *moderate* or *high*
- RCT data starts at *high* and observational data starts at *low*

Vaccine policy recommendations in Australia

- We summarised assessments for cholera, DTP, HPV, influenza, meningococcal, pneumococcal, rabies and zoster vaccines
- Overall, **56%** of policy questions received a GRADE of **low*** or **very low***
- Common reasons were:
  - Potential confounding (e.g., when only observational data was available)
  - Mismatch between the policy question and study-specific research question (e.g., interventions, schedules and populations)
  - Uncertainty in effect estimation (i.e., lack of precision)

- Each disease area had up to 6 research questions
- Research questions had 5-13 outcomes of interest
- Evidence quality varied by question and disease area

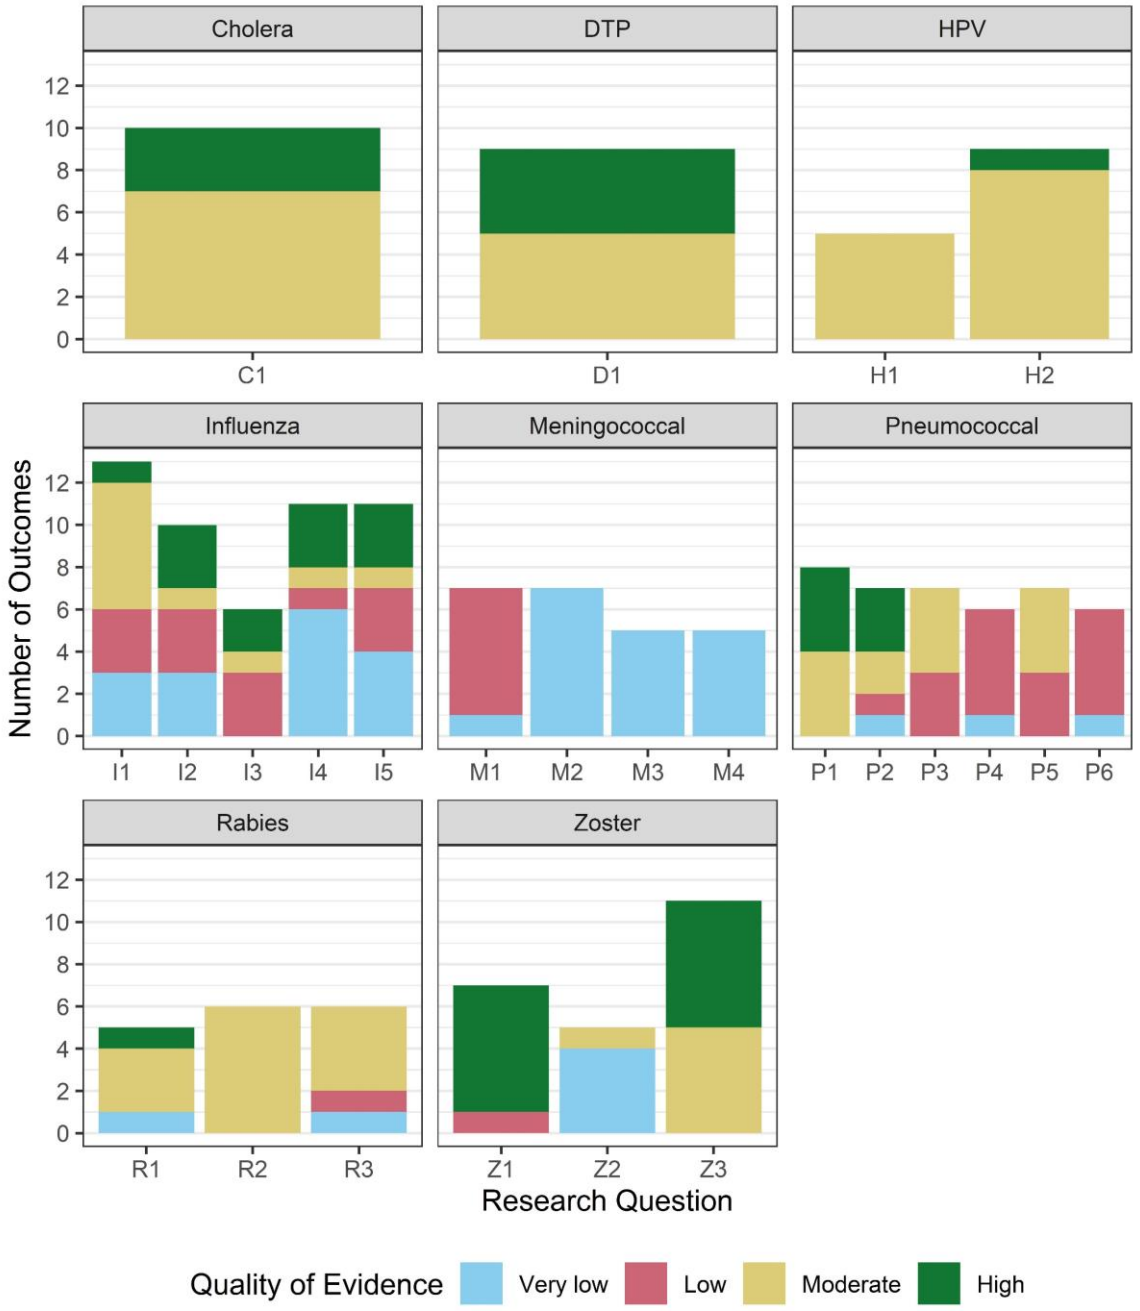

- Outcomes for individual research questions tended to be downgraded for similar reasons

- E.g., pneumococcal and indirectness

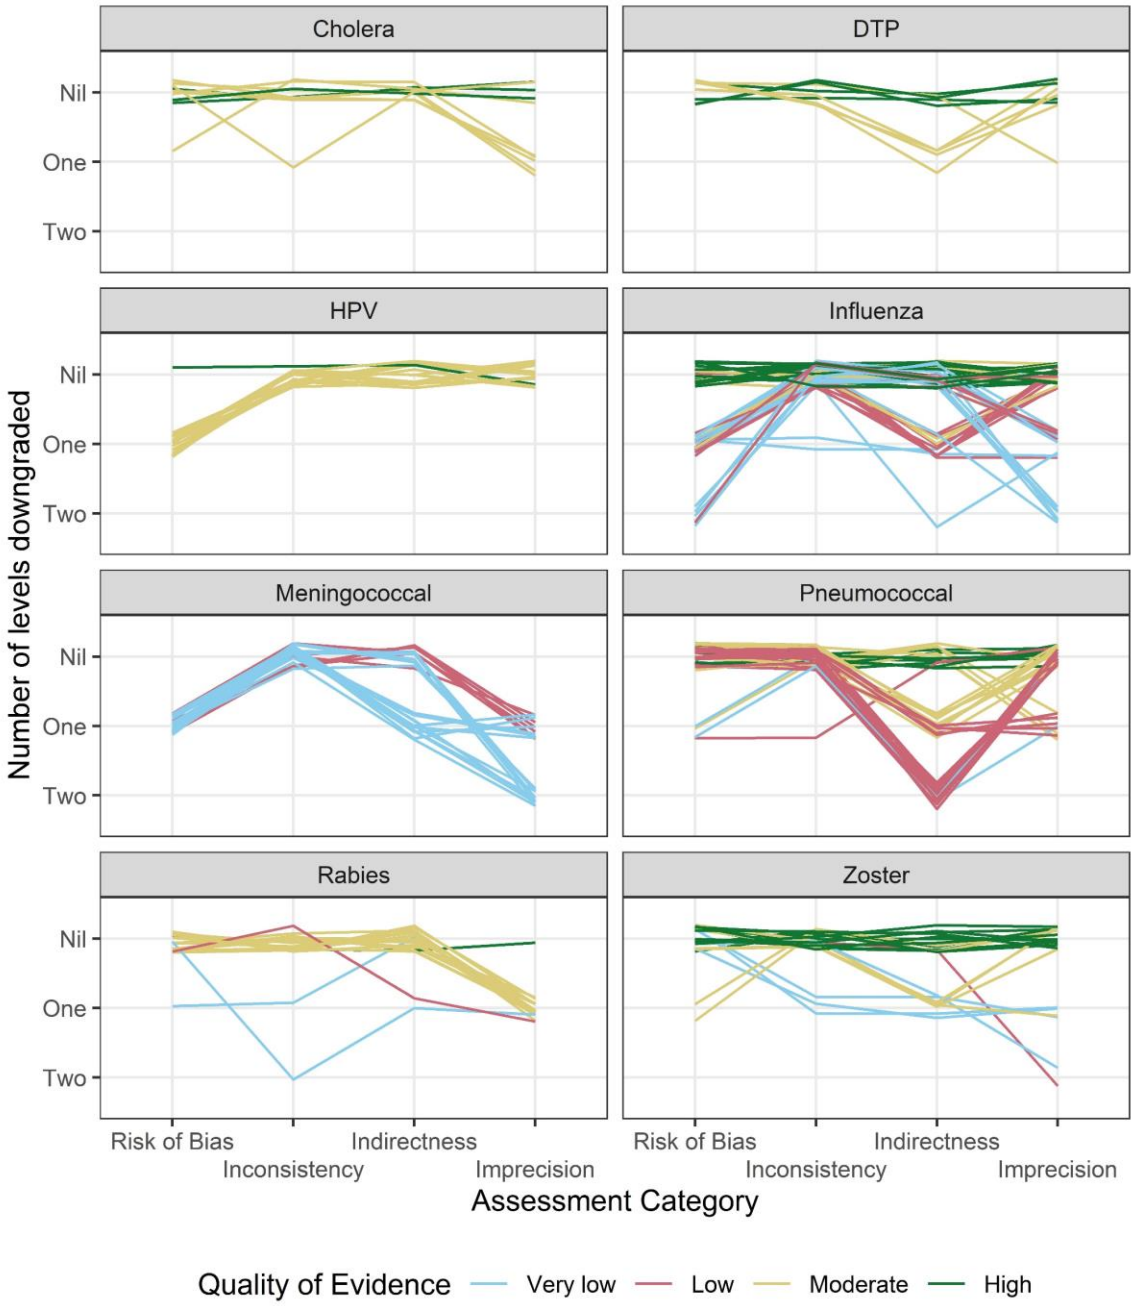

How can we improve the quality of the evidence?

- **Align** trial design with policy question (e.g., interventions and schedules)
- Increase **inclusivity** with respect to targeted subpopulations (e.g., age groups, immunocompromised groups, ethnicity)
- Involve policy-makers in setting research questions **before** the evidence is generated rather than afterwards
- Generate evidence that **informs** the policy-maker’s decision efficiently

A potential solution

The value-driven adaptive design

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 $E_{\boldsymbol{\Theta}}[\mathrm{INB}(\boldsymbol{\Theta})]$

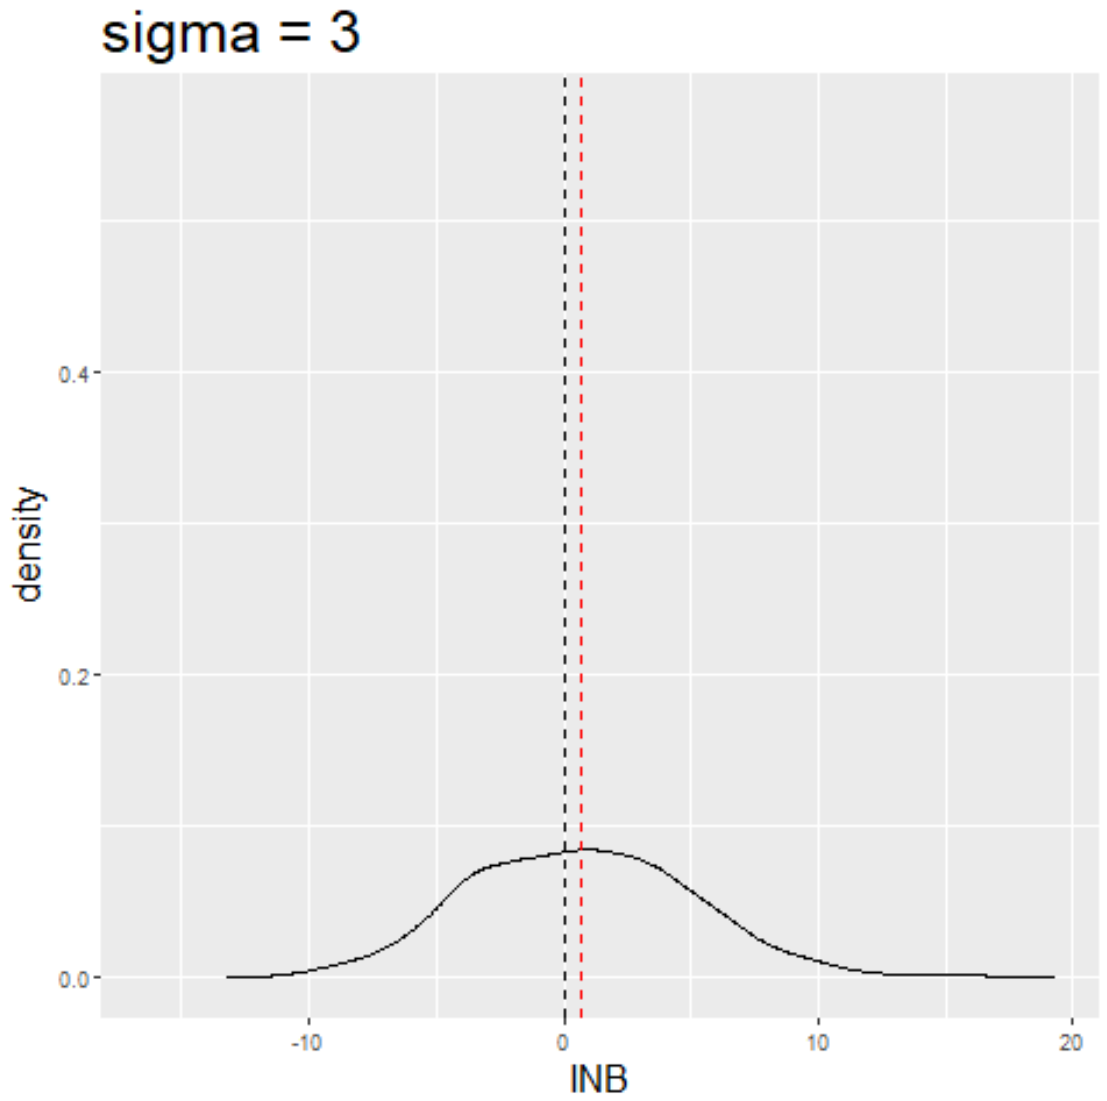

What is the value of reducing uncertainty?

- 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

- I.e., proceed if $\text{EVSI} > \epsilon$, where ϵ is some prespecified cost

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?

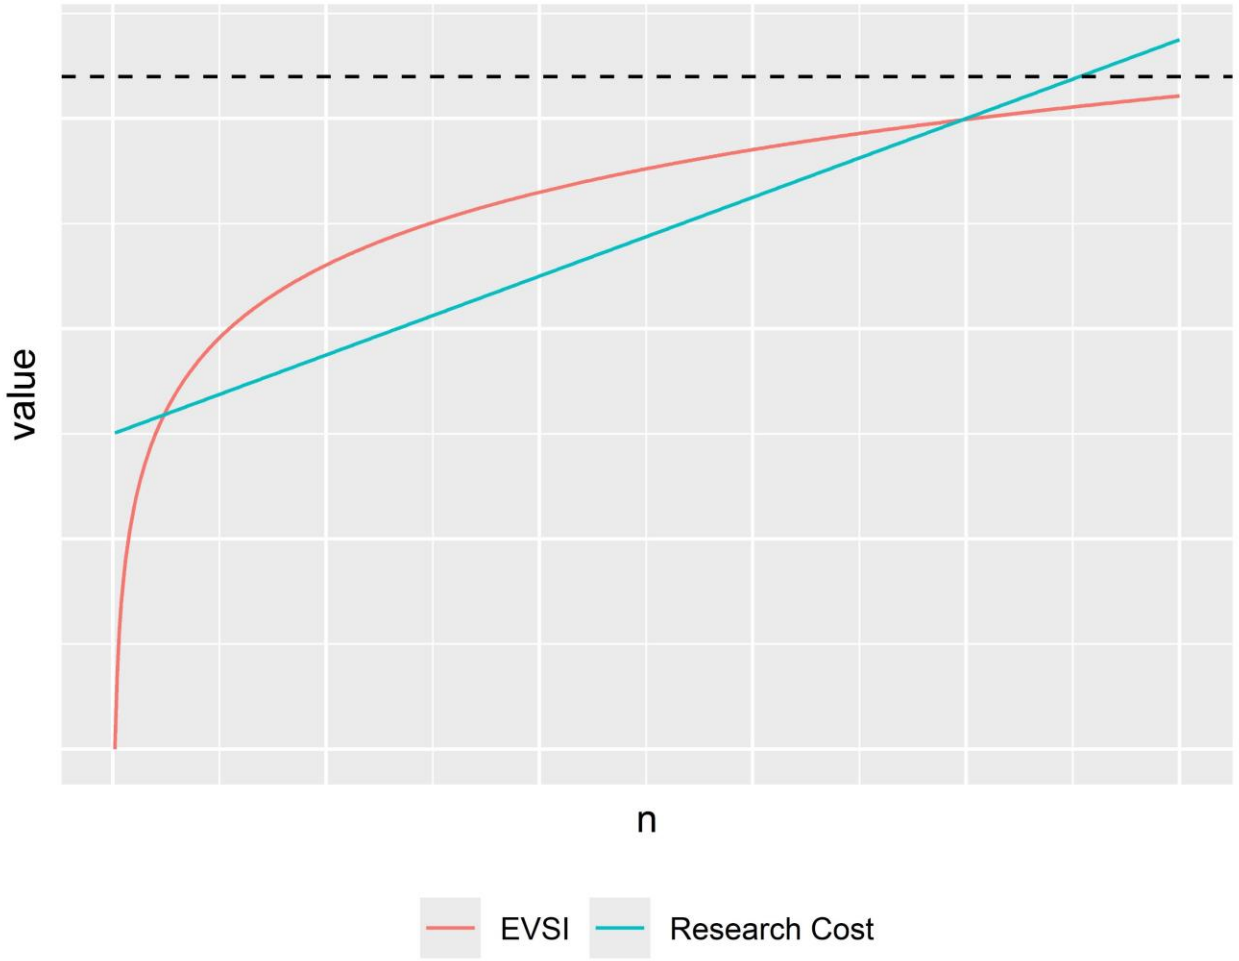

Notation	Definition
$j \in \{1,2,\ldots\}$	Prespecified analyses

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 EVSI at interim analyses to make stop-go decisions
- Suppose the trial has fixed start-up cost δ and per-participant recruitment cost γ
- Start the trial if $\mathrm{EVSI}^0 \geq \delta + \gamma \times n_1$
- Stop at analysis j if $\mathrm{EVSI}^j < \gamma \times n_j$
- Choose the best decision option and report out the results

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 maternal vaccination (MV) or infant immunoprophylaxis (II) will be more cost-effective
- 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

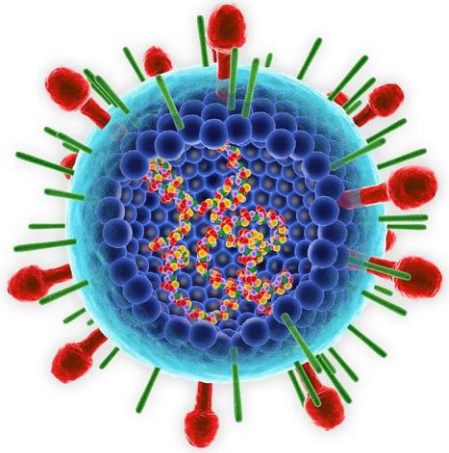

The (simplified) health economic model

MV is more cost-effective than II if:

$$\mathrm{INB}(\boldsymbol{\Theta}) > 0 \leftrightarrow p_{\mathrm{MV}} - p_{\mathrm{II}} < \frac{260}{\lambda}$$

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

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

A clinical trial

- We conduct a clinical trial to reduce the parameter and decision uncertainty
- Randomise up to 1,000 mother-infant dyads (i_k) to each strategy $k \in \{\mathrm{MV}, \mathrm{II}\}$
- Dyads have 12-month MA-RSV outcome $x_{i_k} \in \{0,1\}$
- Analyses after every 250 dyads per strategy $(\mathrm{i.e.,}\ \ n_{1k} = n_{2k} = n_{3k} = n_{4k} = 250)$
- 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: $\sum_{i_k=1}^{\bar{n}_{jk}} X_{i_k} \sim \mathrm{Binomial}(\bar{n}_{jk}, p_k)$

Two illustrative scenarios

1) II is more effective than MV (we set $p_{\mathrm{II}} = 0.10$ and $p_{\mathrm{MV}} = 0.18$)

2) Both strategies are equally effective (we set $p_{\mathrm{II}} = p_{\mathrm{MV}} = 0.10$)

- For both scenarios we estimate $\mathrm{EVSI}^0 \approx \$338\ \mathrm{million}$ using Strong et al.'s non-parametric regression method

- This exceeds the cost of conducting the trial to the first analysis and so we proceed ($\delta + \gamma \times 500 = \$2\ \mathrm{million}$)

Expected value of sample information

Analysis	Scenario 1	Scenario 2
Prior	$338.1 million	$338.1 million
Analysis 1	$5.3 million	$1.0 million
Analysis 2	$0.4 million	<$0.1 million
Analysis 3	<$0.1 million	<$0.1 million
Analysis 4	<$0.1 million	<$0.1 million

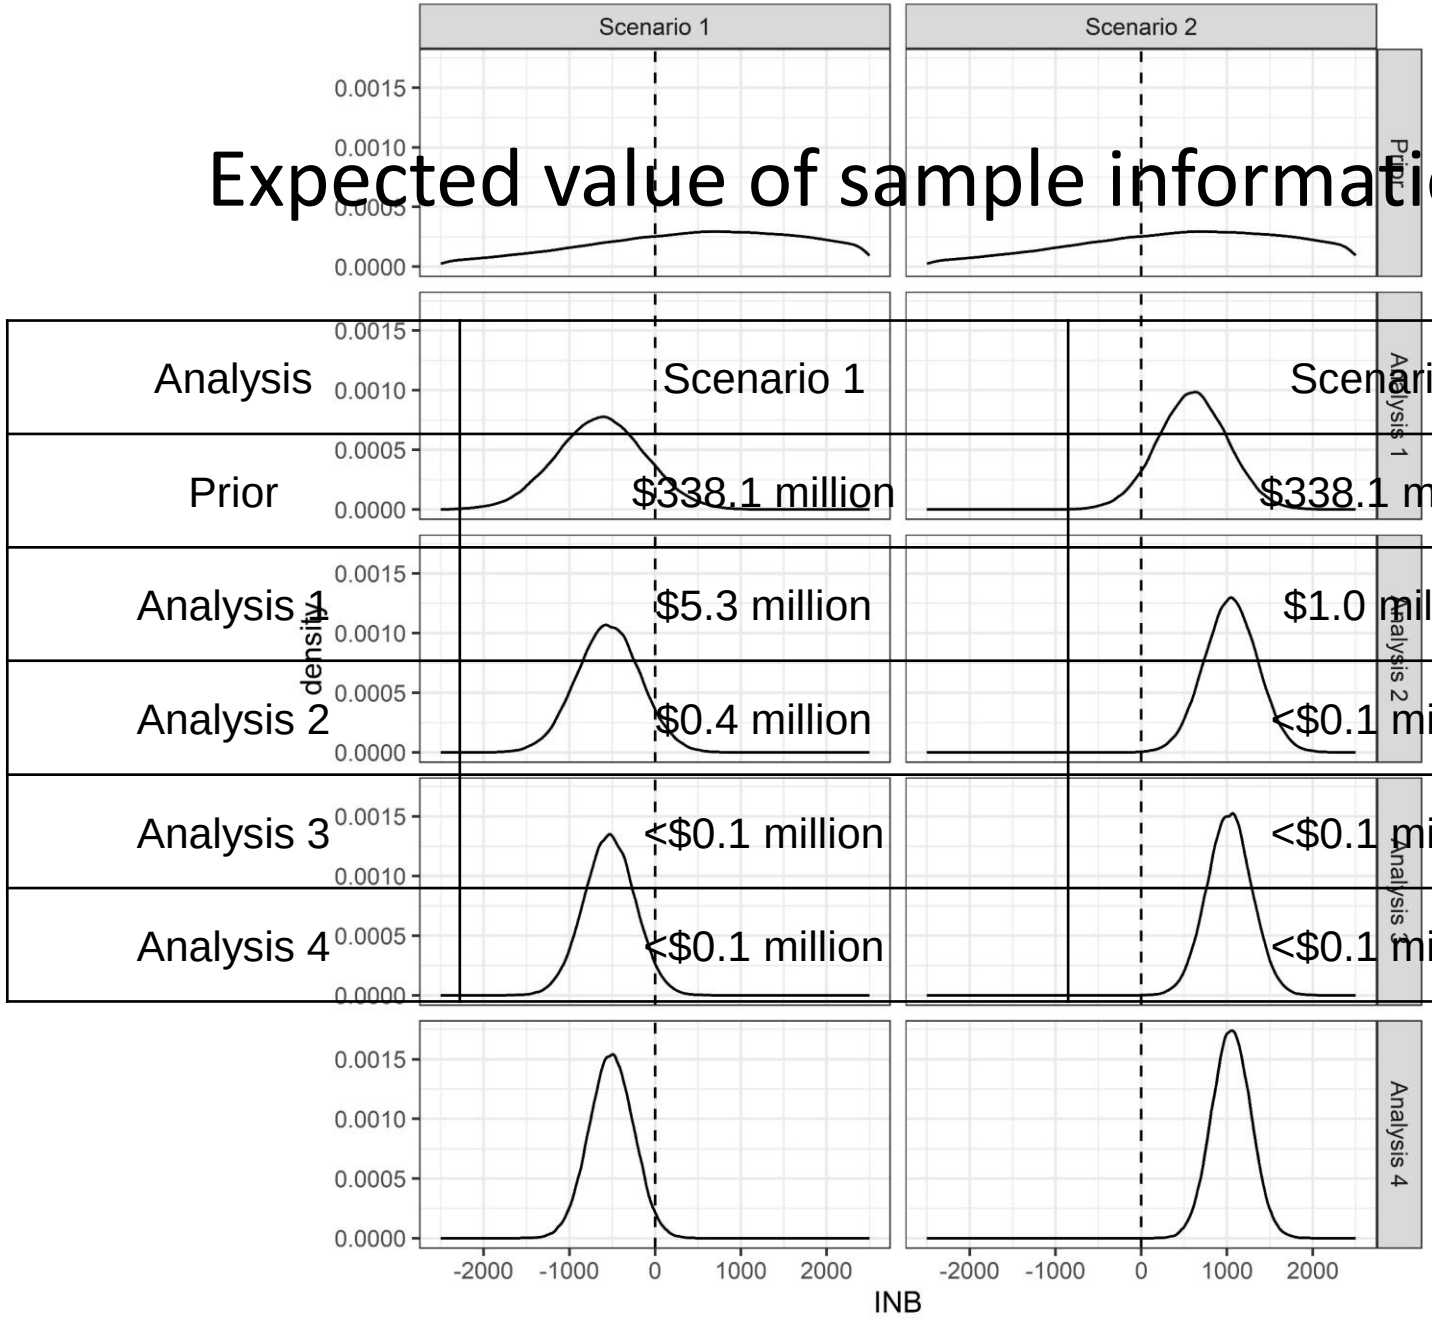

Current state and future directions

- Decision-theoretic methods for optimal selection have been developed over the past few decades
- More recently, Chick, Forster et al. have developed the value-based sequential design for optimal stopping
- Implemented by Flight, Brennan et al. retrospectively to UK funded trials
- Interested in further discussions seeing where our ideas overlap
- One day perhaps we could prospectively design a trial using these methods
- Watch this space

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, incorporation of other decision-theoretic models, etc.)