# Estimating treatment effects from adaptive clinical trials

Ian Marschner

NHMRC Clinical Trials Centre
University of Sydney

ACTA International Clinical Trials Conference, Sydney, October 2019

# Adaptive clinical trials

- The term **adaptive** can mean many things
- **Unifying theme:** Key features of the study design can be changed during the study, so the design can **adapt** to information accumulated during the conduct of the study
- Adaptive clinical trials may have advantages in terms of:
  - flexibility
  - speed
  - resource utilisation
  - ethics
- **In this talk we will focus on estimating treatment effects**

# Adaptive design features

- Number of participants
  - Early stopping (benefit, harm, futility); sample size re-estimation
- Randomisation ratio between treatments
  - Adaptive randomisation (play-the-winner, Bayesian adaptive *etc.*)
- Number of stages
  - Single arm multi-stage phase II (Simon design, Gehan design *etc.*)
- Number of treatments
  - Multi-arm multi-stage designs (MAMS)
- Treatment doses
  - Adaptive dose-finding; up-and-down designs; random walk designs
- Types of participants
  - Population enrichment designs; subgroup selection

# Response-adaptive designs

- In response-adaptive designs the design features adapt on the basis of unblinded treatment responses
- Particularly complex from a statistical point of view
- **Fundamental statistical issue:**
  - Unlike a standard fixed (non-adaptive) study, the study design in a response-adaptive clinical trial is *random* and it contains information about the treatment effect
- **Implication:**
  - Standard treatment effect estimates based on fixed designs may exhibit bias when used in adaptive clinical trials
- **Analysis of adaptive clinical trials should involve an assessment of whether the treatment effect estimate is likely to be biased**

# Bias in treatment effect estimates

- Treatment effect bias in adaptive clinical trials has two types

- **Unconditional bias:**
  - Average difference between estimated effect and population effect, averaged over all study designs that *could* occur
  - Typically small. Hence, meta-analyses of adaptive studies lead to unbiased treatment effect estimation.

- **Conditional bias:**
  - Average difference between estimated effect and population effect, for the particular study design that *did* occur
  - Can be large or small. Hence, individual adaptive studies may lead to underestimation, overestimation, or unbiased estimation of effects.

# Example 1: stopping early for benefit

- Fan, DeMets and Lan. *J Biopharm Stat* 2004;14:505–530.
- Conditional and unconditional bias in sequential clinical trials

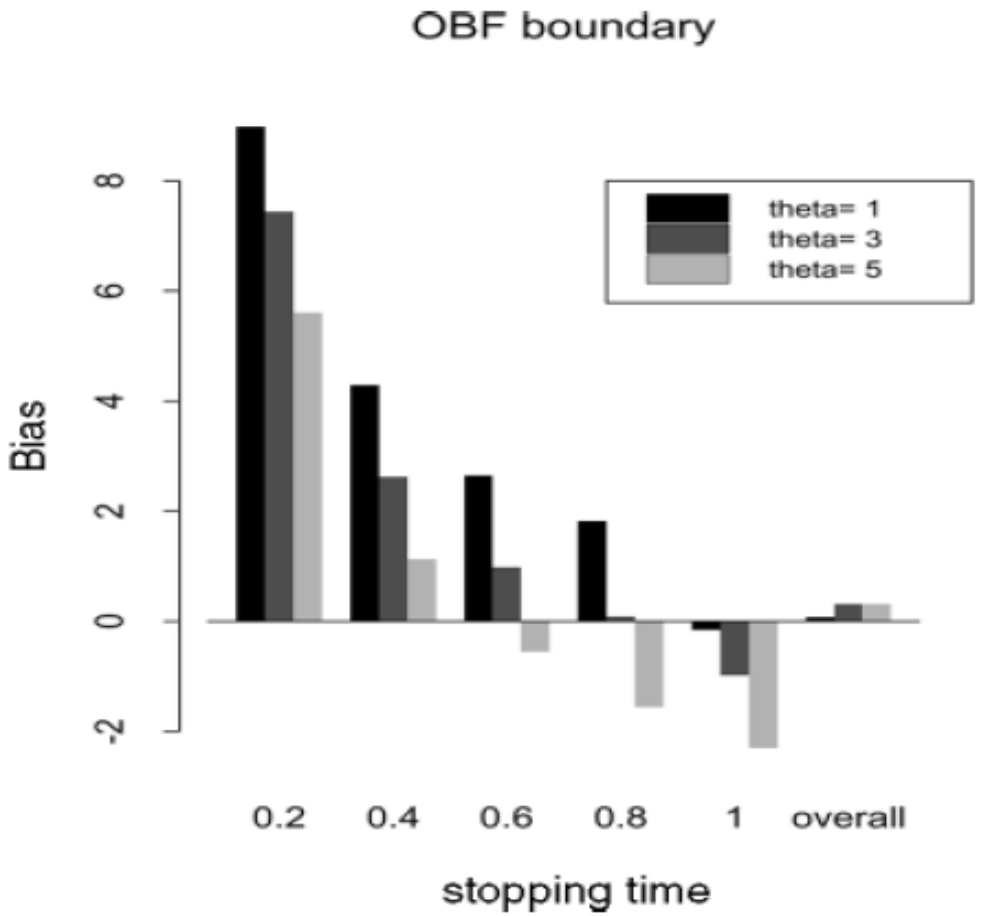

# Example 2: adaptive randomisation

- Randomised play-the-winner design: 10,000 simulations
- $\Pr(\text{response}) = 0.5$ (experimental) or $0.25$ (control)
- Average experimental allocation $= 60\%$
- Actual allocation varies randomly

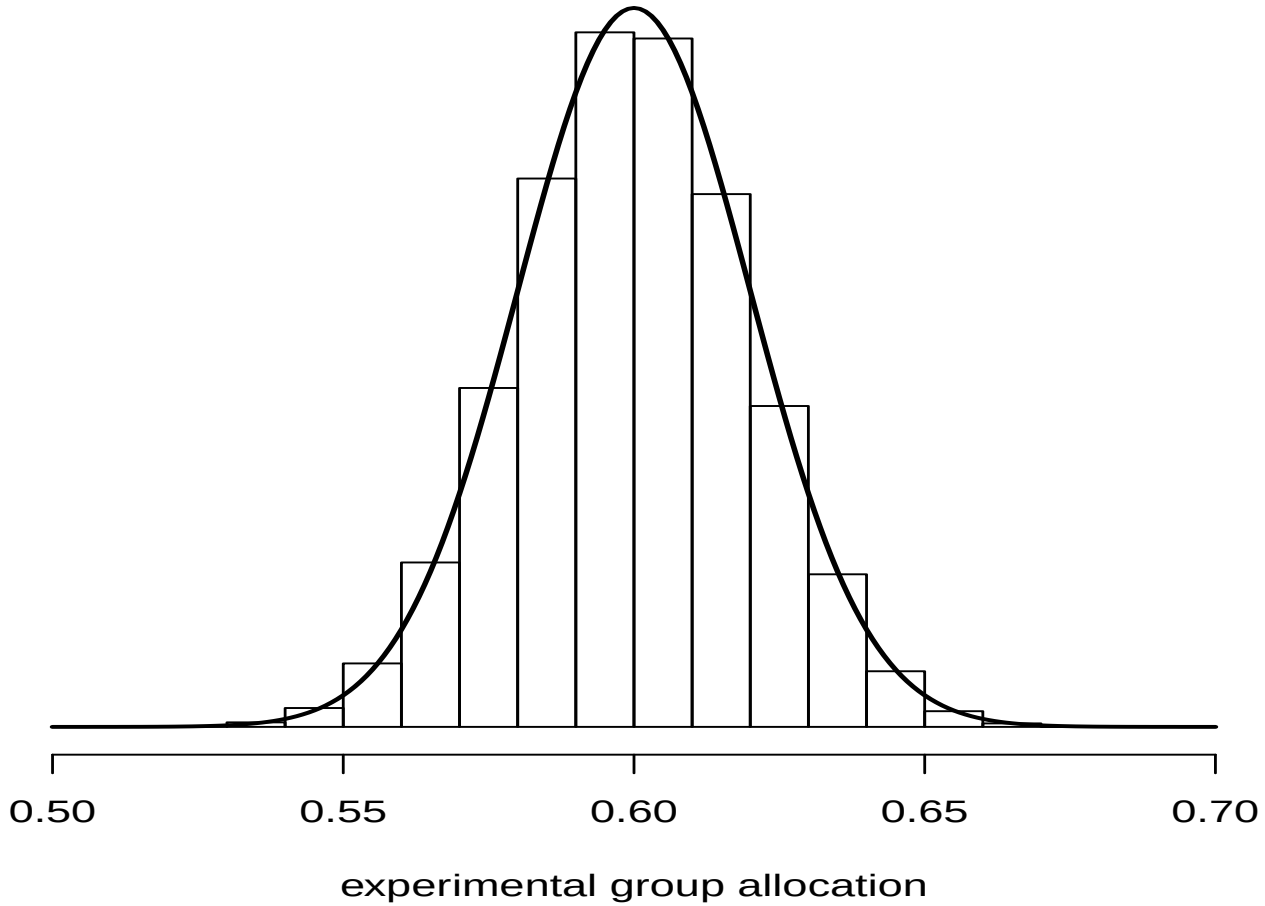

# Example 2: adaptive randomisation

- When allocation is close to 60% treatment effect is unbiased
- When allocation is not close to 60% treatment effect is biased
- Averaged over all possible allocations treatment effect is unbiased

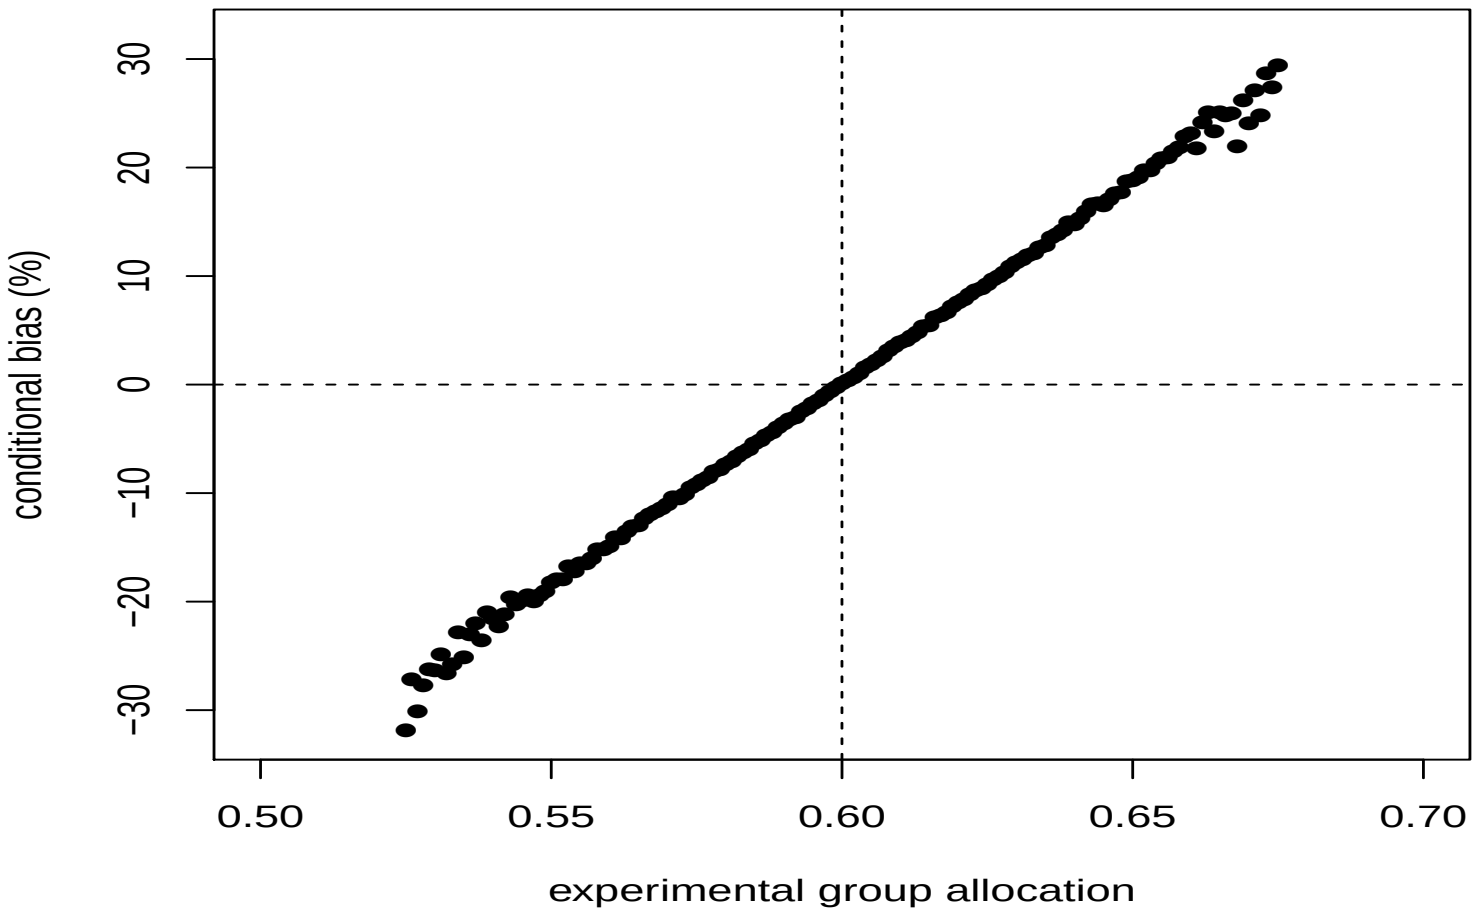

# Example 3: multi-stage phase II

- 3-stage extension of Simon Design
- Stage 1: $n = 9$ Stage 2: $n = 13$ Stage 3: $n = 23$

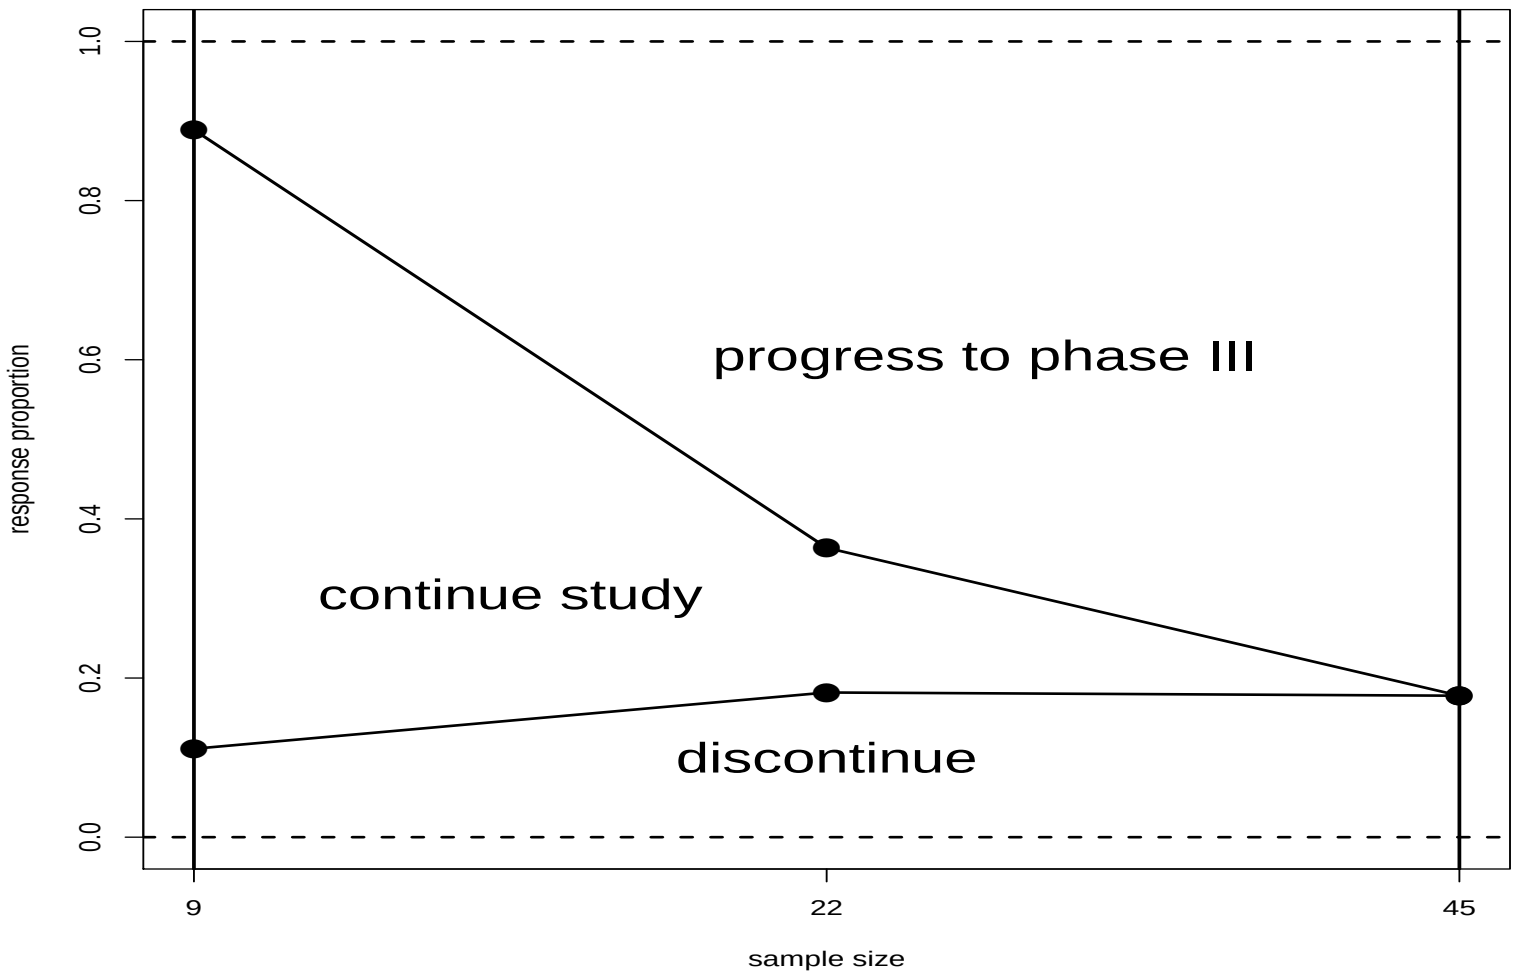

# Example 3: multi-stage phase II

- 10,000 simulations of studies with various population response rates
- Average observed response rate in studies that reached stage 3

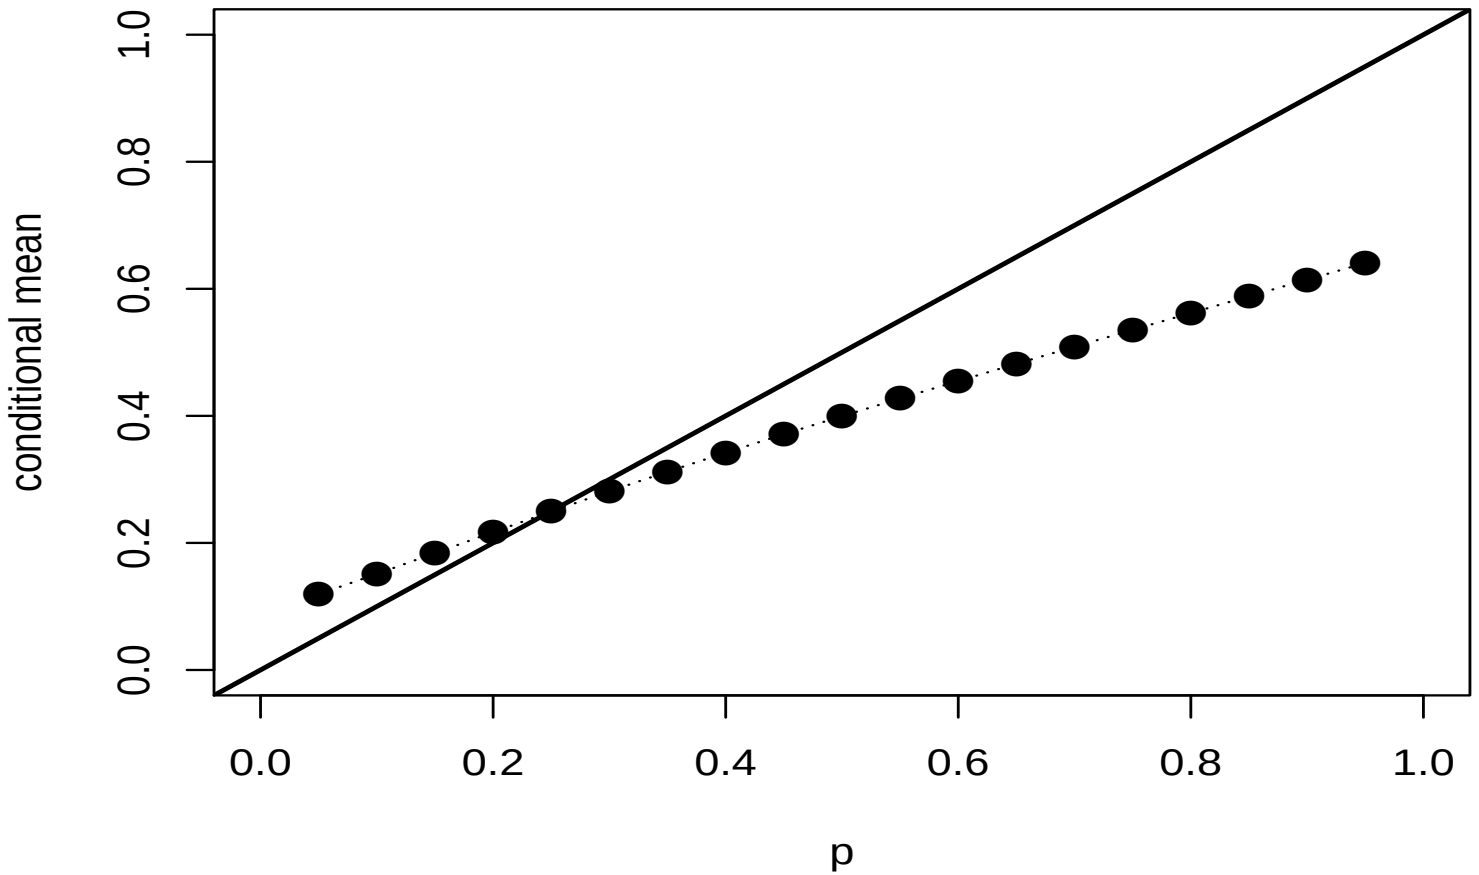

# Conditional statistical inference

- A general analysis method that interprets the observed treatment effect using knowledge of the specific design that *did* occur, rather than with reference to all designs that *could* have occurred
- The fact that a particular design configuration has occurred, gives us information about how the observed treatment effect will behave
- *e.g.* If a study stops for benefit at the first interim analysis, we know the observed treatment effect estimate will be an overestimate
- By using knowledge of how treatment effects tend to behave for particular design configurations, we can adjust for conditional bias
- **Conditional maximum likelihood estimation** can be used in any adaptive study to remove the conditional bias of standard analyses

# Example 2: adaptive randomisation

- Conditional MLE for randomised play-the-winner simulations
- Unbiased given any observed allocation ratio

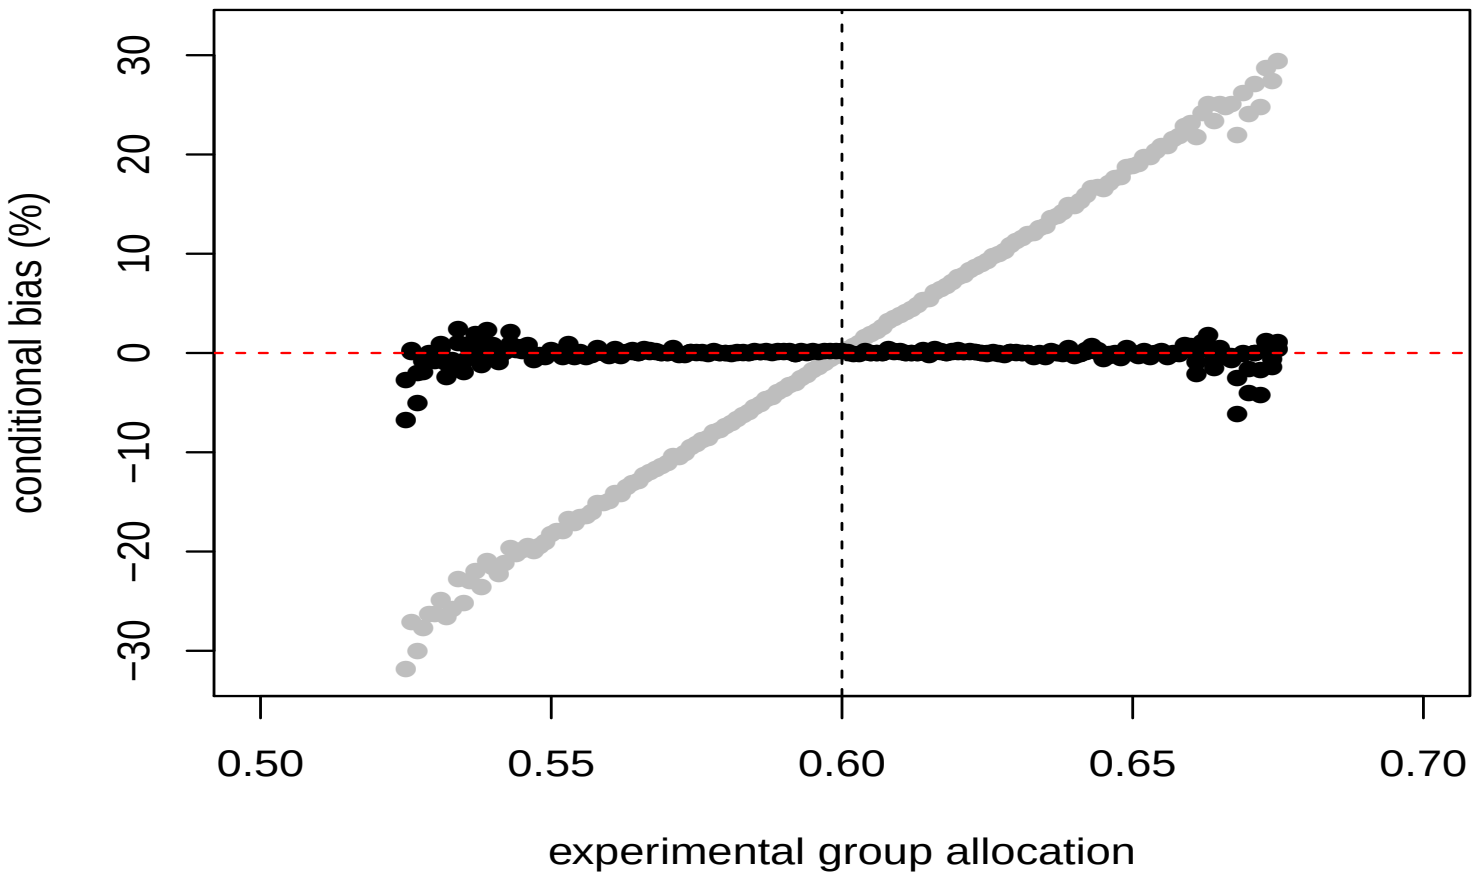

# Example 3: multi-stage phase II

- Conditional MLE for multi-stage study that goes to final stage
- Unbiased regardless of the underlying population response rate

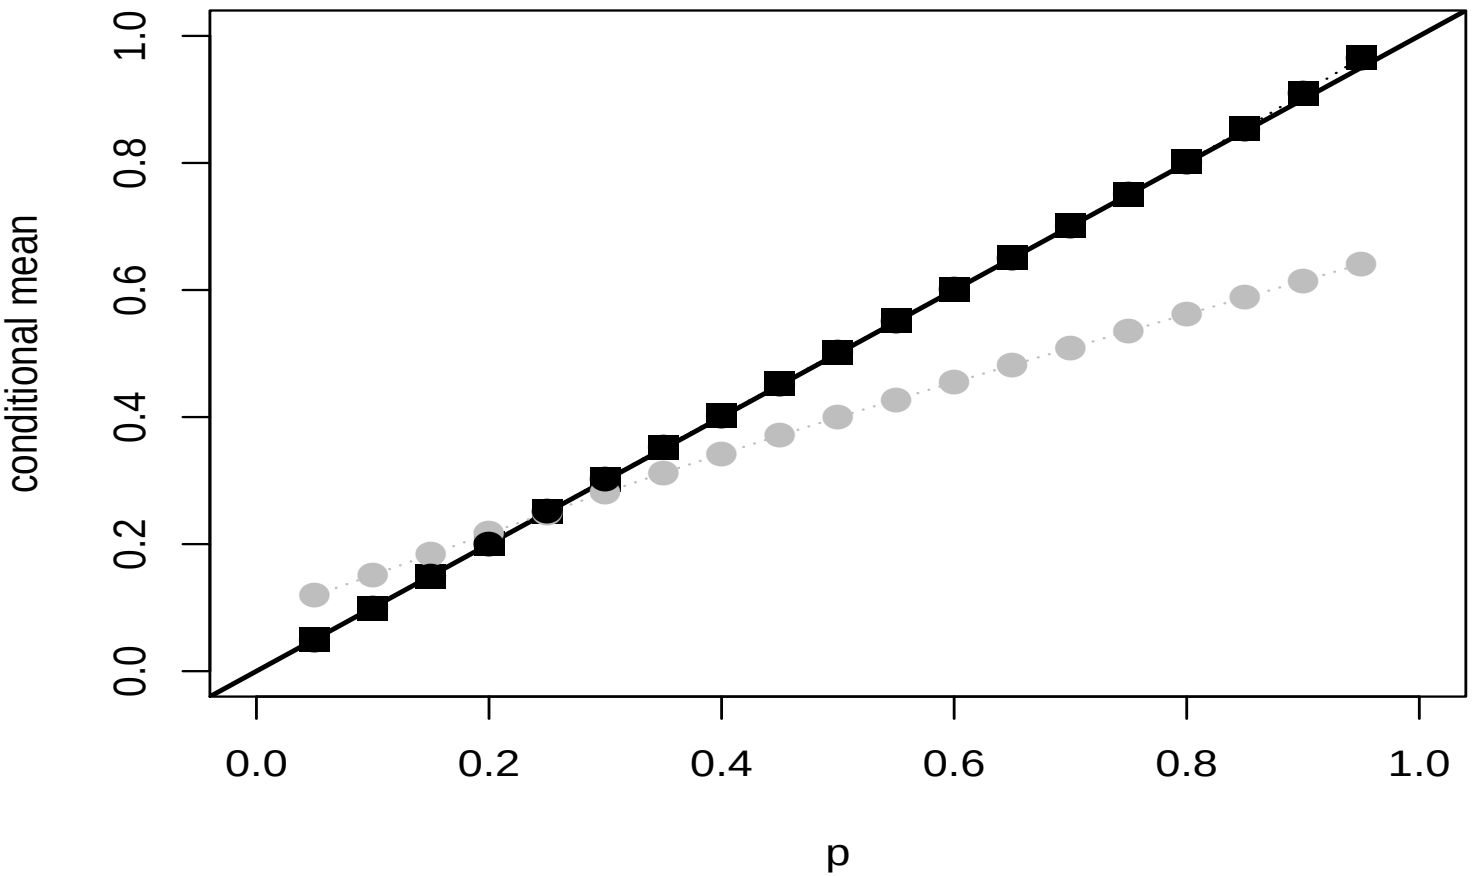

# No free lunch

- Although unbiased, some information may be lost
- This may mean greater variability (or less statistical efficiency) for observed design configurations that are not subject to bias

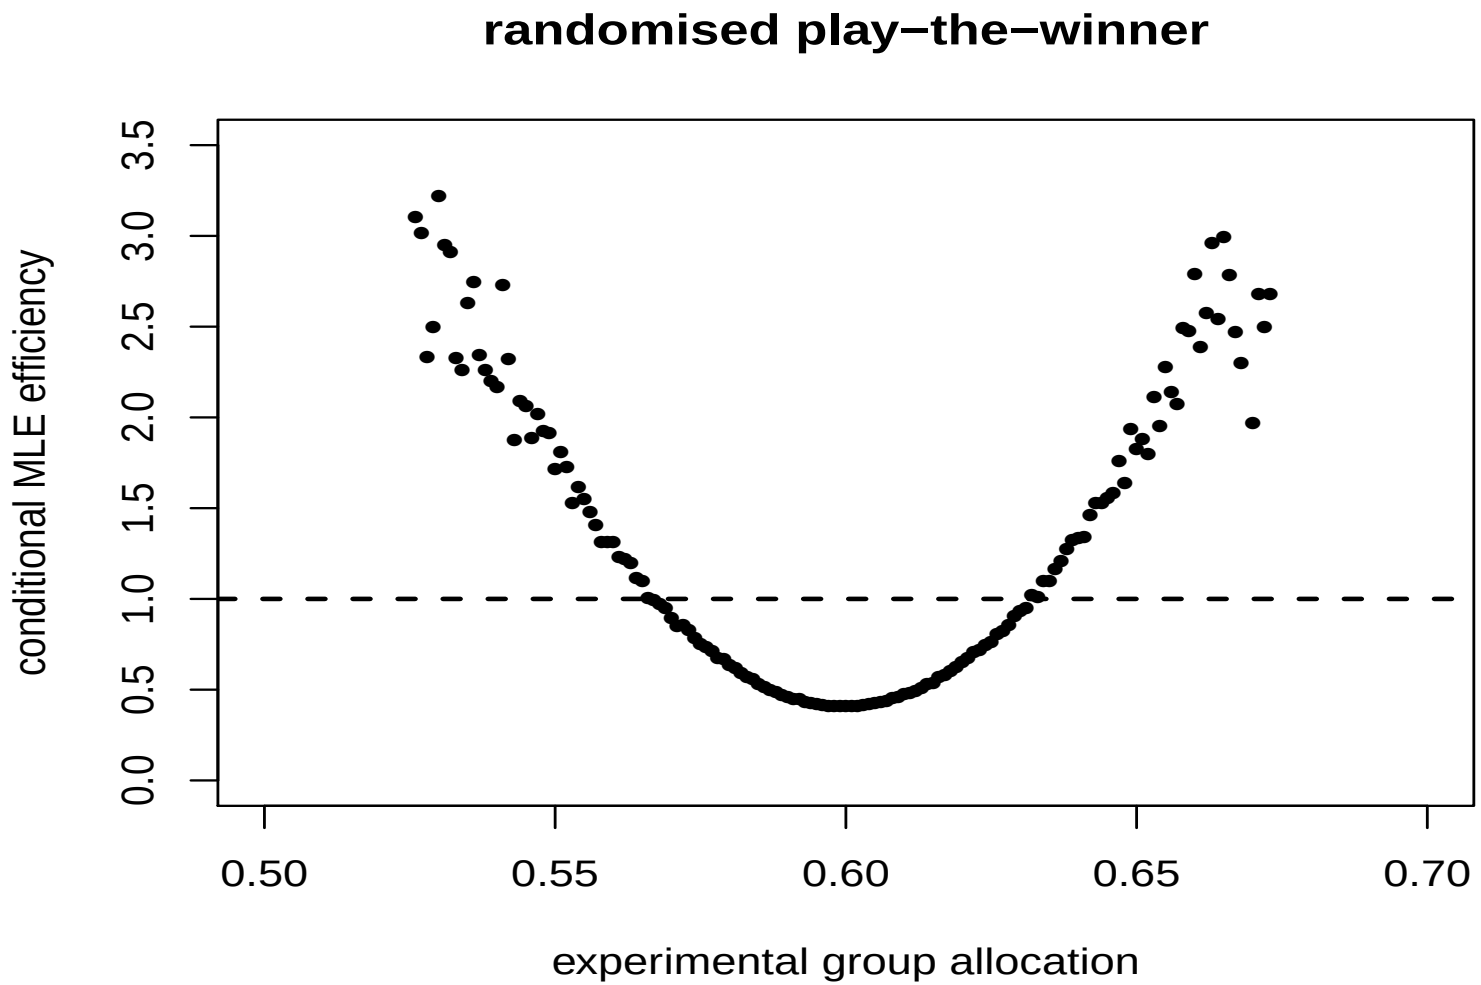

# Conclusions

- Adaptive clinical trials have a study design that is random and is potentially informative about the treatment effect
- For a given realisation of an adaptive design, a standard analysis method may produce a biased treatment effect estimate
- Methods to adjust for this potential bias are available but are not always as statistically efficient as standard methods
- **Analyses of adaptive clinical trials should examine the potential for estimation bias in the treatment effect estimate**
- **Conditional estimation procedures are a useful supplement to standard analysis methods for adaptive clinical trials**