

# Power Analysis & Sample Size Calculation

(Require Substance-Matter Knowledge)

Last modified: 2026-01-09

## Overview

- Large effects from subtle manipulations?
- Inference and power
- Power analysis by simulation
- Do it yourself

Help, my effect size is too large!

# Help, my effect size is too large!

- Probably never heard anyone complain about this
- But it is a huge problem for the scientific integrity of our field
- Reported effect sizes in the literature are way too large

# Help, my effect size is too large!

- Probably never heard anyone complain about this
- But it is a huge problem for the scientific integrity of our field
- Reported effect sizes in the literature are way too large
- Examples
  - Decision biases from two-hand tapping
  - Beautiful parents have more daughters

## Refresher: Framing

- Tversky and Kahneman (1981)

“Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed” (p. 453)

If Program A is adopted **200** people will be **saved** [109]

If Program B is adopted there is 1/3 probability that **600** people will be **saved**, and 2/3 probability that **no people** will be **saved** [43]

If Program C is adopted **400** people will **die** [34]

If Program D is adopted there is 1/3 probability that **nobody** will **die**, and 2/3 probability that **600** people will **die** [121]

- Odds ratio (OR) = 9.0

## Decision biases from two-hand tapping

- McElroy and Seta (2004),  $n = 48$

“a behavioral task of finger tapping was used to induce asymmetrical activation of the respective hemispheres . . . Framing effects were found when the right hemisphere was selectively activated whereas they were not observed when the left hemisphere was selectively activated” (p. 572)

|      | right-hand tapping |       | left-hand tapping |       | ratio of odds ratios (ROR) |
|------|--------------------|-------|-------------------|-------|----------------------------|
|      | safe               | risky | safe              | risky |                            |
| gain | 8                  | 4     | 12                | 1     |                            |
| loss | 7                  | 4     | 3                 | 9     |                            |
| OR   |                    | 1.1   |                   | 36    | 31.5                       |

- Our replication (see Gelman, 2020),  $n = 332$

|      |    |     |    |     |     |
|------|----|-----|----|-----|-----|
| gain | 52 | 31  | 56 | 27  |     |
| loss | 26 | 57  | 30 | 53  |     |
| OR   |    | 3.7 |    | 3.7 | 1.0 |

# Beautiful parents have more daughters

- Kanazawa (2007)

“Very attractive individuals are 26% less likely to have a son” (p. 133)

–  $n_{total} = 2970$   
–  $n_{v.att.} < 400$

- Gelman and Weakliem (2009)

“the noise is stronger than the signal” (p. 314)

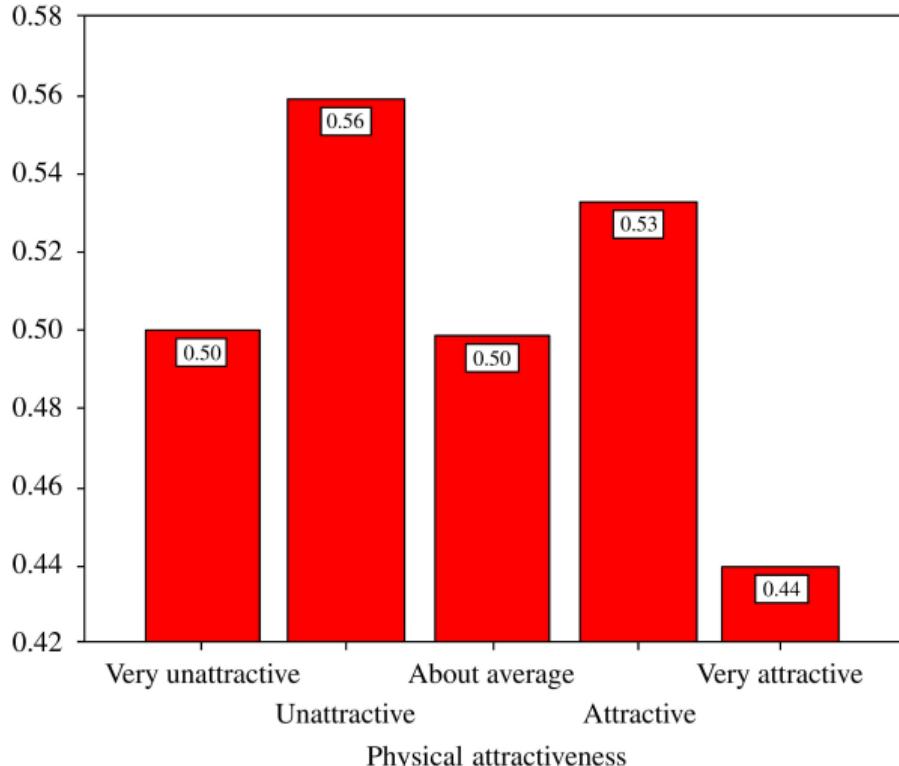
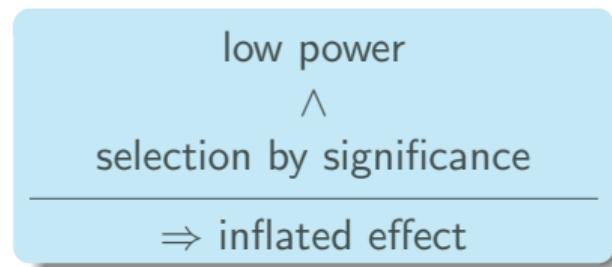


Fig. 1. Proportion of boys among the first child, by parent's physical attractiveness.

## Large effects from subtle manipulations?

There is a simple explanation for the seemingly large effects published all over the psychological literature

- that works without any real large effects
- but assumes that they are statistical artifacts based on a combination of



(type M error; Gelman & Carlin, 2014)

## Classical inference in a nutshell

- Deciding between two hypotheses about parameter of data-generating model (Neyman & Pearson, 1933)
- Null hypothesis (specific), alternative hypothesis (logical opposite)
  - Example: Binomial model,  $H_0: \pi = 0.5$ ,  $H_1: \pi \neq 0.5$
- Possible decision errors

|            | Decision for $H_0$     | Decision for $H_1$     |
|------------|------------------------|------------------------|
| $H_0$ true | correct                | type I error, $\alpha$ |
| $H_1$ true | type II error, $\beta$ | correct                |

### Conventions

- $\alpha = 0.05$
- $\beta < 0.2$

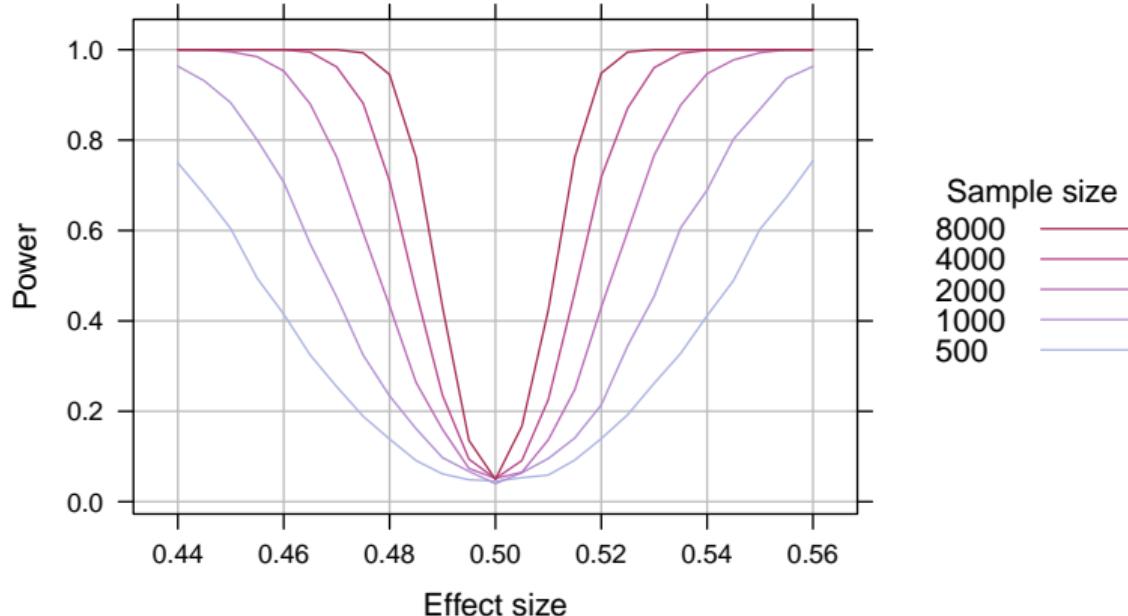
- Decision based on data (p-value)
  - If  $p < \alpha$ , choose  $H_1$ ; else retain  $H_0$
- Power =  $1 - \beta$ 
  - Probability of test to detect an effect of a given size

# Power function

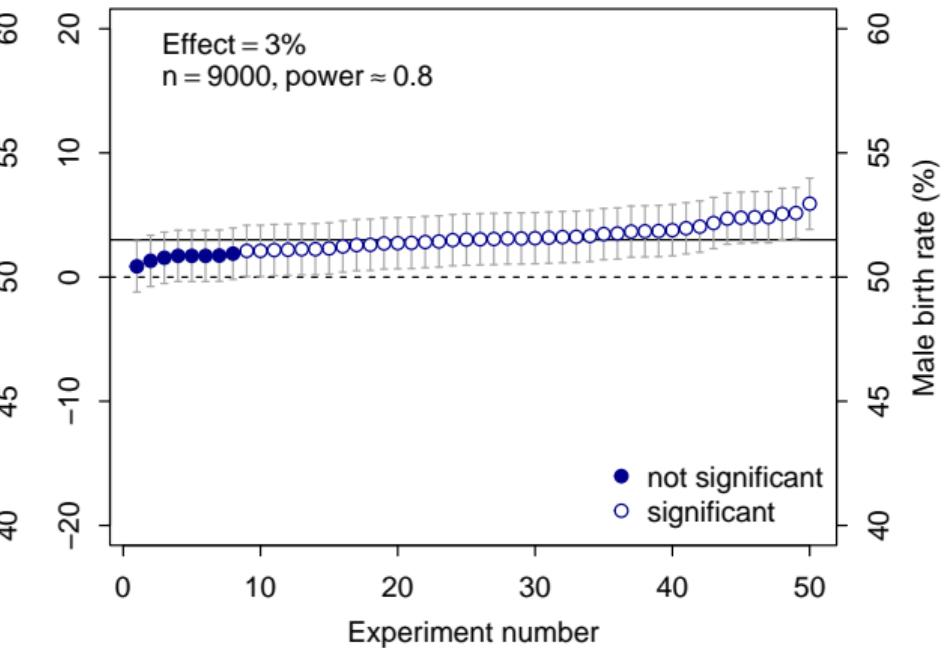
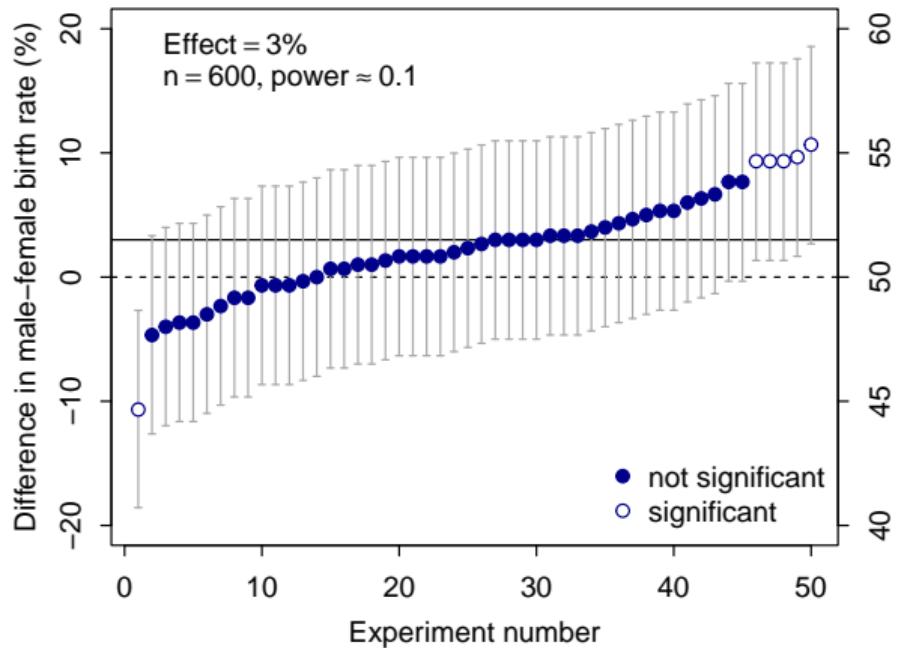
Power of a test depends on

- effect size  
(deviation from  $H_0$ )
- sample size  $n$
- $\alpha$

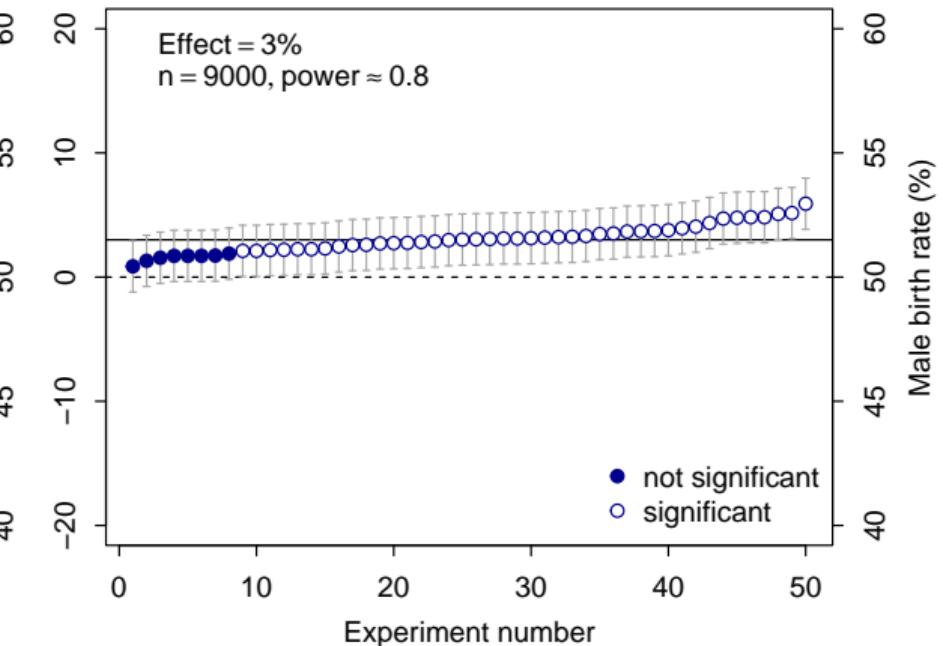
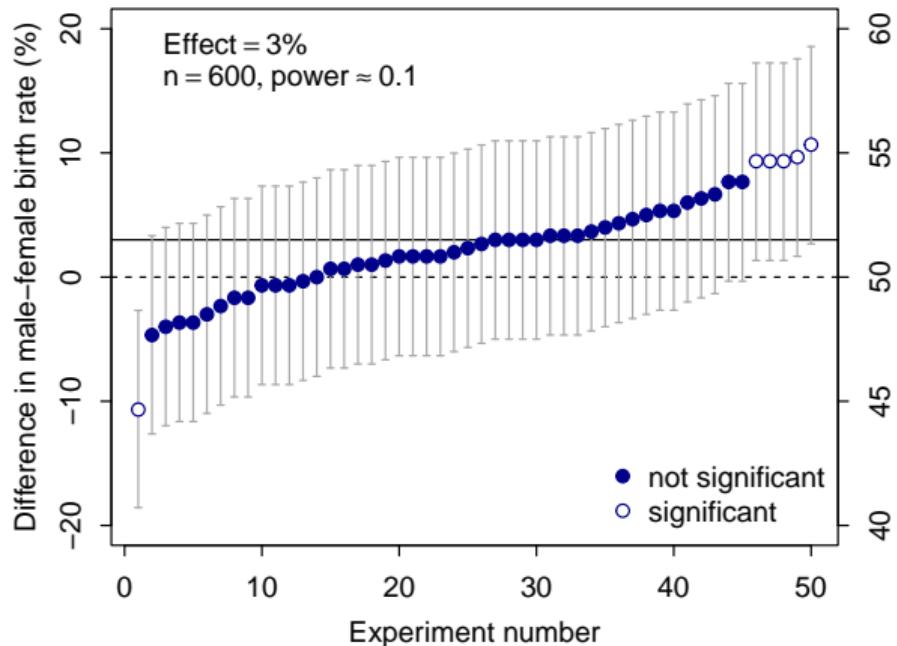
With effect size, power, and  $\alpha$  fixed, we can calculate  $n$



# High power is a necessary condition for valid inference



# High power is a necessary condition for valid inference



“If power is low . . . every possible outcome under repeated sampling will be misleading: there will be a high proportion of inconclusive null results, and any significant effects will be due to mis-estimations of the true effect” (Vasishth & Gelman, 2021, p. 1317)

## Exercise: First steps in simulation

- Generate data from a binomial model using the `rbinom()` function in R; try out different values of
  - $n$  (10, 500, 2000)
  - the parameter  $\pi$  (0.5, 0.8, 0.44, 0.515)and see how this affects the output
- With these data, test different null hypotheses using `binom.test()`; these may or may not coincide with the values of  $\pi$  used for data generation
- If you repeat data generation and testing, can you usually reject  $H_0$ ?

## Power analysis by simulation

### Why simulation?

- Simulation is at the heart of statistical inference
- Inference: Compare the data with the output of a statistical model
- If data look different from model output, reject model (or its assumptions)
- Simulation forces us to **specify a data model** and to attach meaning to its components
- Model should not be totally unrealistic for those aspects of the world we want to learn about

## Power simulation

The steps in general

1. Specify the model including the effect of interest
2. Generate observations from the model
3. Test  $H_0$
4. Repeat

Power is estimated from the proportion of significant test results

## Specify the model including the effect of interest

(1) Choose statistical model according to its assumptions

- Binomial test → binomial distribution → `rbinom()`
- t-test → normal distribution → `rnorm()`
- ...

## Specify the model including the effect of interest

(1) Choose statistical model according to its assumptions

- Binomial test → binomial distribution → `rbinom()`
- t-test → normal distribution → `rnorm()`
- ...

(2) Fix unknown quantities

- Standard deviations, correlations, ...
- Plausible values from the literature (beware of significance filter)

## Specify the model including the effect of interest

(1) Choose statistical model according to its assumptions

- Binomial test → binomial distribution → `rbinom()`
- t-test → normal distribution → `rnorm()`
- ...

(2) Fix unknown quantities

- Standard deviations, correlations, ...
- Plausible values from the literature (beware of significance filter)

(3) Specify the effect of interest

- *Not* the true effect (else no need to run the study!)
- *Not* the effect one expects or hopes to find (size of effect is unknown!)
- *Never* an effect size taken from another study (significance filter!)
- *But* the biologically or clinically or psychologically “relevant effect one would regret missing” (Harrell, 2020)

## Power simulation and sample size

The steps in pseudo code

```
1  Set sample size to n
2  replicate
3  {
4      Draw sample from model with minimum relevant effect
5      Test null hypothesis
6  }
7  Determine proportion of significant results
```

Sample size calculation

- Sample size  $n$ , minimal relevant effect and  $\alpha$  must be predetermined
- Adjust  $n$  until desired power (0.8 or 0.95) is reached
- To be on the safe side, assume higher variation, less (or more) correlation, and smaller interesting effects (what results can we expect, if ... )

## Power simulation example

### Binomial test

- Birth rates in the general population



51.5 %



48.5 %

→ corresponds to a male-female sex ratio of 106:100

### More examples

- Wickelmaier (2022) includes power simulation examples and R code for many classical statistical tests

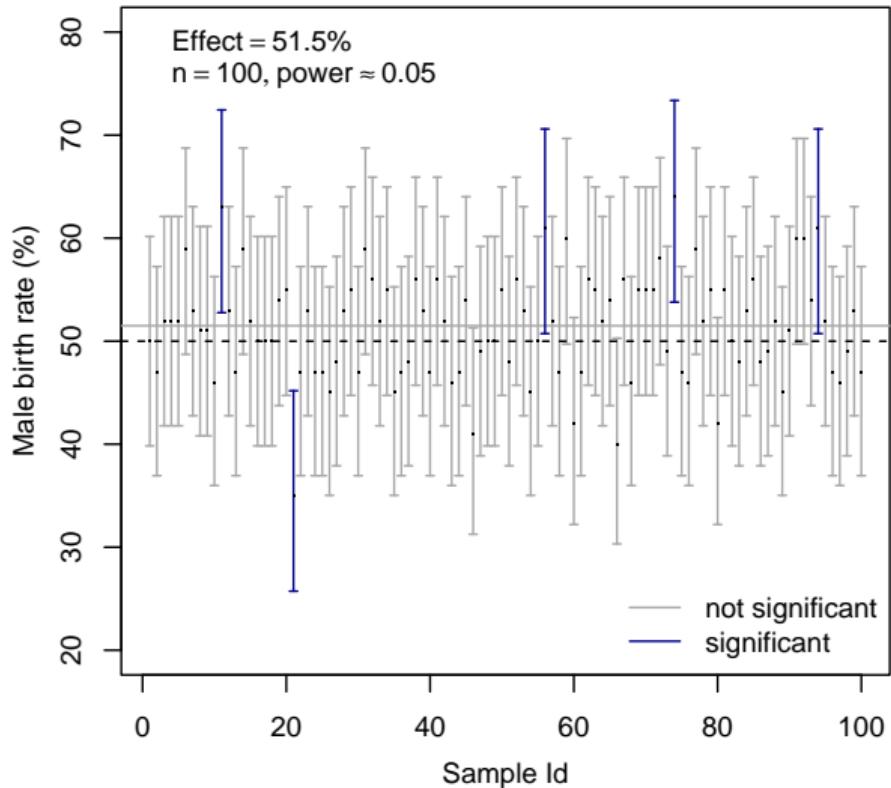
## Example: Birth rates

- Fisher's principle states that the male-female sex ratio is about 1:1
- Plan a study and calculate the sample size necessary to
  - detect a deviation from Fisher's principle of 106:100
  - with about 80% power
- Check your setup
  - Set the effect size to zero; what “power” estimate do you expect to get?

```
1 n <- ...                                # adjust sample size
2 pval <- replicate(5000, {                # replications of experiment
3   x <- rbinom(1, size = n,               # data-generating model with
4       prob = 106/(106 + 100))           # minimum relevant effect
5   binom.test(x, n = n, p = 1/2)$p.value # p-value of test against H0
6 })
7 mean(pval < 0.05)                         # simulated power at alpha = 0.05
```

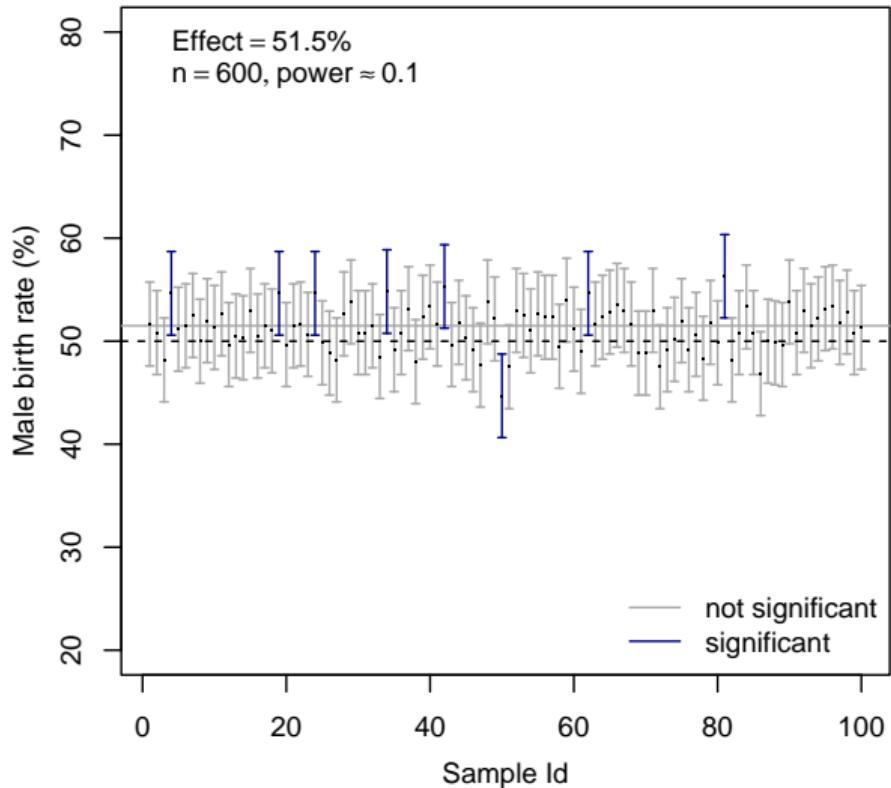
## Example: Birth rates

Fisher's principle: Testing against  $H_0: \pi = 0.5$



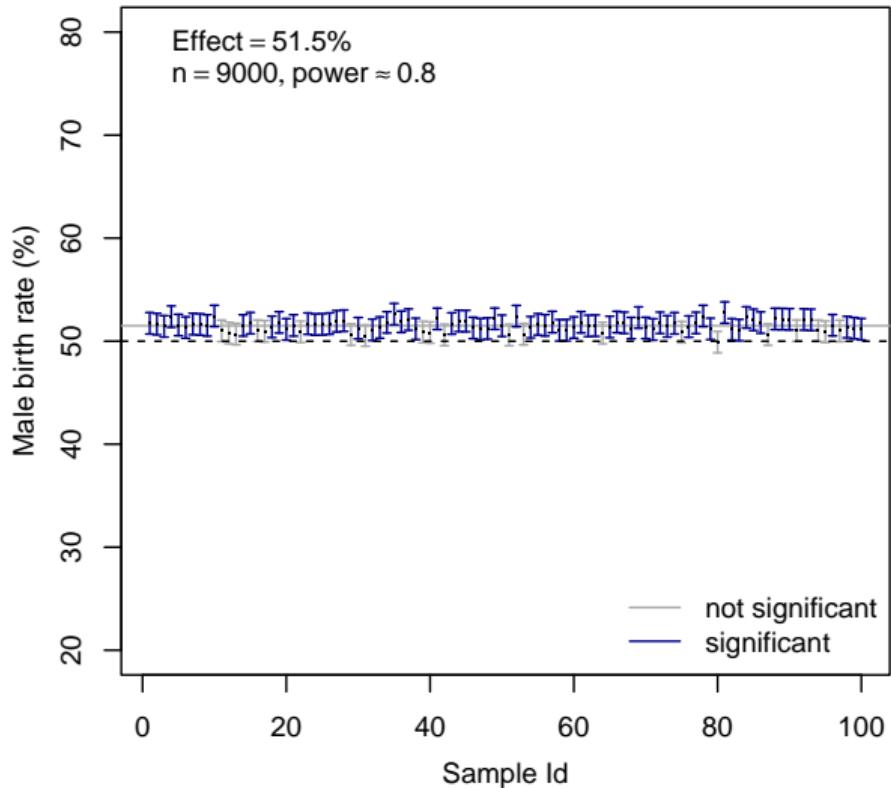
## Example: Birth rates

Fisher's principle: Testing against  $H_0: \pi = 0.5$



## Example: Birth rates

Fisher's principle: Testing against  $H_0: \pi = 0.5$



## Final thoughts

Statistical tests are no screening procedures

- Significance is not a substitute for relevance
- Nonsignificance does not imply absence of effect

- Often, data are rather uninformative and compatible with many models and hypotheses
- At the same time, “all models are wrong” (Box, 1976)
- Making data-based decisions using statistical inference requires a confirmatory setting where a-priori substantive knowledge goes into the power analysis
- When relying on statistical tests outside such a setting, all we do is descriptive statistics with p-values; this does more harm than good

## References I

Box, G. E. P. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356), 791–799. doi: 10.1080/01621459.1976.10480949

Gelman, A. (2020, October 22). *An odds ratio of 30, which they (sensibly) don't believe*. Retrieved from <https://statmodeling.stat.columbia.edu/2020/10/22/an-odds-ratio-of-30-which-they-sensibly-dont-believe/>

Gelman, A., & Carlin, J. (2014). Beyond power calculations: Assessing type S (sign) and type M (magnitude) errors. *Perspectives on Psychological Science*, 9(6), 641–651. doi: 10.1177/1745691614551642

Gelman, A., & Weakliem, D. (2009). Of beauty, sex and power: Too little attention has been paid to the statistical challenges in estimating small effects. *American Scientist*, 97(4), 310–316. doi: 10.1511/2009.79.310

Harrell, F. (2020, June 20). *Statistical problems to document and to avoid*. Retrieved from <https://discourse.datamethods.org/t/author-checklist/>

Kanazawa, S. (2007). Beautiful parents have more daughters: A further implication of the generalized Trivers–Willard hypothesis (gTWH). *Journal of Theoretical Biology*, 244(1), 133–140. doi: 10.1016/j.jtbi.2006.07.017

McElroy, T., & Seta, J. J. (2004). On the other hand am I rational? Hemispheric activation and the framing effect. *Brain and Cognition*, 55(3), 572–580. doi: 10.1016/j.bandc.2004.04.002

## References II

Neyman, J., & Pearson, E. S. (1933). On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London, Series A*, 231(694–706), 289–337.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453–458. doi: 10.1126/science.7455683

Vasishth, S., & Gelman, A. (2021). How to embrace variation and accept uncertainty in linguistic and psycholinguistic data analysis. *Linguistics*, 59(5), 1311–1342. doi: 10.1515/ling-2019-0051

Wickelmaier, F. (2022). Simulating the power of statistical tests: A collection of R examples. *ArXiv*. doi: 10.48550/arXiv.2110.09836

## P-value

The p-value is the probability of obtaining a test statistic that signals a deviation from  $H_0$  at least as extreme as that observed in the experiment, given  $H_0$  is true and its underlying model holds

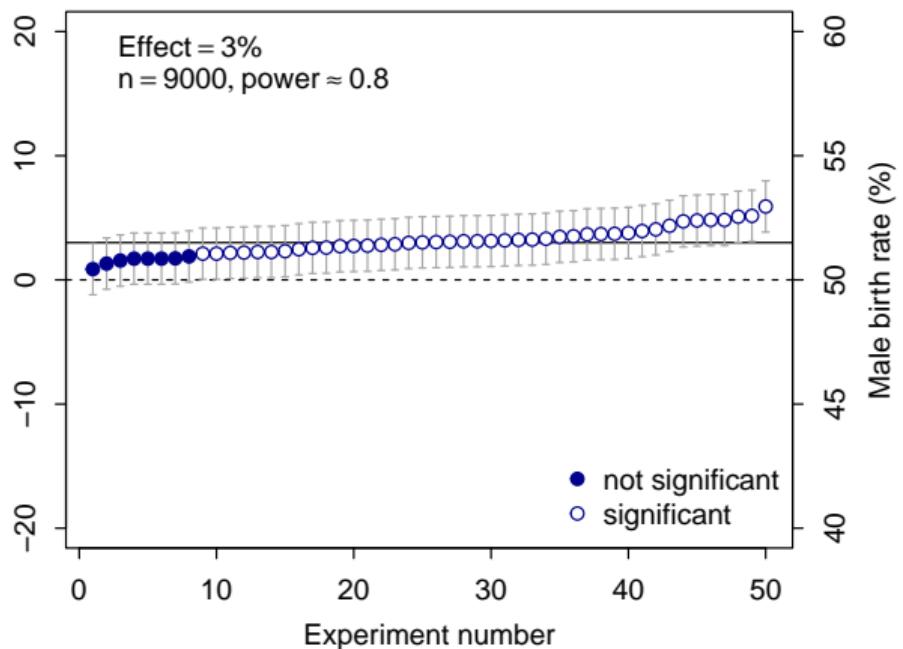
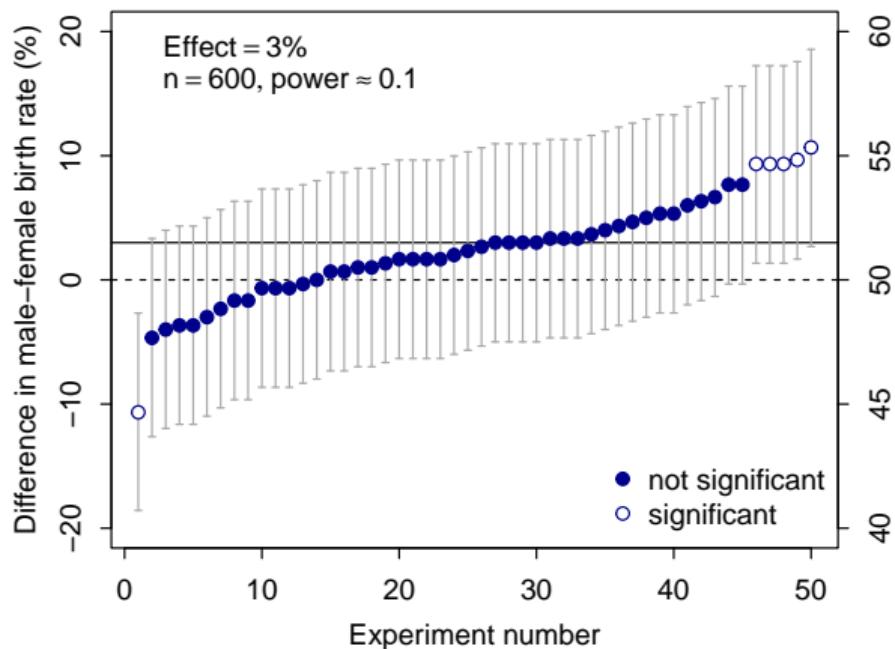
<https://apps.mathpsy.uni-tuebingen.de/fw/pvalbinom/>

## On the role of power

- Vasisht and Gelman (2021)

“the importance of power cannot be stressed enough. Power should be seen as the ball in a ball game; it is only a very small part of the sport, because there are many other important components. But the players would look pretty foolish if they arrive to play on the playing field without the ball. Of course, power is not the only thing to consider in an experiment; no amount of power will help if the design is confounded or introduces a bias in some way” (p. 1333)

# Binomial test power simulation



<https://apps.mathpsy.uni-tuebingen.de/fw/birthrate/>