Considerations of sample size and power calculations given a range of analytical scenarios

Short title: Considerations of sample size and power calculations in research

Alice R Carter^{1,2}, Kate Tilling^{1,2,3}, Marcus R Munafò^{1,4}

1. Medical Research Council Integrative Epidemiology Unit, University of Bristol, Bristol, UK

2. Population Health Sciences, Bristol Medical School, University of Bristol, Bristol, UK

3. National Institute for Health Research Bristol Biomedical Research Centre, University Hospitals Bristol NHS Foundation Trust and University of Bristol.

4. School of Psychological Science, University of Bristol, Bristol, UK

Corresponding author:

Dr Alice R Carter

Oakfield House,

Oakfield Grove,

Bristol,

BS8 2BN

alice.carter@bristol.ac.uk

Contributors

All authors contributed to the study design, manuscript preparation and final draft. KT and MRM

conceived the original systematic review project which informed this manuscript, assisted in

interpreting the results and made comments on and assisted in drafting the manuscript. ARC carried

out all analyses, reviewed manuscripts included in the systematic review and wrote the manuscript.

Funding

ARC, KT and MRM all work in a unit which receives core funding from the Medical Research Council

and University of Bristol (MC_UU_00011/1, MC_UU_00011/3 and MC_UU_00011/7), which

additionally funds ARC. This study was supported by the NIHR Biomedical Research Centre at

University Hospitals Bristol NHS Foundation Trust and the University of Bristol. The views expressed

28 in this publication are those of the author(s) and not necessarily those of the NHS, the National
29 Institute for Health Research or the Department of Health. No funders were involved in this project.

30 Conflicts of interest

31 All authors no conflicts of interest.

32

Abstract

The sample size of a study is a key design and planning consideration. However, sample size and power calculations are often either poorly reported or not reported at all, which suggests they may not form a routine part of study planning. Inadequate understanding of sample size and statistical power can result in poor quality studies. Journals increasingly require a justification of sample size, for example through the use of reporting checklists. However, for meaningful improvements in research quality to be made, researchers need to consider sample size and power at the design stage of a study, rather than at the publication stage. Here we briefly illustrate sample size and statistical power in the context of different research questions and how they should be viewed as a critical design consideration.

43 Introduction

44 The replicability of published scientific studies is the subject of considerable ongoing debate (1-5).
45 Adequate sample size is a key design consideration to ensure a study can answer the research
46 question of interest, where sample size planning can be used to either achieve adequate statistical
47 power, or accurate estimates of a parameter (or both) (6). Statistical power is the probability of
48 rejecting the null hypothesis at a specified significance level, given a true effect of a specified
49 magnitude. In a null hypothesis significance testing (NHST) framework, a smaller sample size (all
50 other things being equal) results in lower statistical power for a given effect size. In an estimation
51 framework, lower sample size reduces the precision of estimates. The impact of low power is
52 exacerbated by publication bias, where positive (i.e. significant) findings are more likely to be
53 published than null (i.e. non-significant) findings (7). In studies with low statistical power to detect
54 even moderate effects, significant effects are likely to be an over-estimate of the true effect;
55 therefore, published studies are likely to be biased compared with the true effect of an independent
56 variable on a dependent variable (1, 8).

57 Statistical power is a central element of the NHST approach used widely across the biomedical
58 sciences, where the inference is simply reduced to whether the null hypothesis is rejected or not, at
59 a specified significance level, alpha, which is typically set to 5%. Although some have argued that this
60 conventional level should be lower (9). Power depends on several factors: sample size, the minimum
61 effect size of interest, the variability of the outcome, and the significance (alpha) level. Some study
62 designs have more power for the same number of participants than others, by reducing the
63 variability of the outcome. For example, a within-subjects design has greater power to detect a given
64 effect size, compared with an otherwise identical between-subjects design for the same number of
65 participants, because within-person variability is usually less than between-person variability (10). In
66 general, the larger the effect size of interest, the smaller the sample size needed to attain a given
67 power to detect it. Whilst a power of 80% is conventionally considered adequate, there is evidence

from a broad range of biomedical disciplines that the average power of published studies may be much lower (11, 12).

Sample size is also important when estimating effects, with a greater or lesser degree of precision (e.g., as reflected in confidence intervals), rather than simply testing a null hypothesis. Here, sample size considerations focus on obtaining sufficiently precise confidence intervals for an effect estimate, (13). These methods are explained with formulae by Rothman and Greenland (13).

Given the growing literature on the prevalence of inadequate sample sizes, and the likely impact on the robustness of many reported findings, journals are becoming increasingly aware of the need to act (8, 12). One measure adopted by some journals is asking authors to complete a reporting checklist (14). These require authors to report a range of details about their methodology and statistical analyses alongside their submission, often including how sample sizes were decided (14-16). However, in the absence of a good understanding of why sample size and power calculations are important (particularly at the design stage) and how these should be conducted in different contexts, the value of these checklists may be limited.

We reviewed whether the introduction of checklists is associated with more authors reporting how they determined their sample size (see Box 1). The introduction of checklists coincided with an increase in the inclusion of some mention of sample size, compared with articles submitted to the same journal before the checklists were introduced and with journals where checklists had not been introduced over the same period. However, there were *fewer* formal sample size or power calculations reported. In studies that commented on their sample size, justification was largely given in terms of no statistical methods or tests were used or that the sample was similar to the field (17). This suggests a proper consideration of sample size considerations may be lacking, particularly at the study design stage.

Box 1: Empirical example of reporting checklists

Reporting checklists have been introduced by some journals (e.g., *Nature Publishing Group* journals) to improve research transparency and the reporting of statistical methods. As part of this, authors must now describe how they determined their sample size for their study, and why this is adequate. However, in a small systematic review we carried out to investigate the effects of these reporting checklists we identified some potential unintended consequences of these checklists (17).

In our review, comparing 36 studies submitted in the 3 months prior to reporting checklists being introduced with 45 studies submitted after checklists were introduced, we saw a *reduction* in formal power calculations being carried out. Although the number of studies reporting some kind of justification of their sample size in the article main text increased after checklists were implemented from 25% (95% confidence interval [CI]: 12% to 42%) to 53% (95% CI: 38% to 68%), the number of formal power calculations decreased from 14% (95% CI: 9% to 21%) to 9% (95% CI: 2% to 21%).

In the post-intervention group, 24 studies (53% [95% CI: 38% to 68%]) commented on how their sample sizes were determined in their main text, but 20 (44% [95% CI: 30% to 60%]) of the studies stated that no statistical methods (i.e., a formal power calculation) were used or that the sample sizes were similar to those typical for the field (which, if the average power of previous studies in a field is low, as has been shown in many cases, will simply perpetuate low power).

This suggests that while the checklist intervention did have the intended consequence of ensuring authors improved the transparency of their approaches and methods, there was little improvement in the methods used to determine sample size, meaning that any improvements in research reproducibility and reliability may be limited. Since the number of statements on sample size increased, this suggests inadequate justification of study sample size has also increased.

Here we present a brief overview of different research scenarios and appropriate approaches to deciding on the appropriate sample size. We have developed a flowchart (Figure 1) following four analytical scenarios outlining when and how to estimate sample size or power, and how this should impact on the rest of the planned study. Firstly, we consider two scenarios when a new study is being set up and primary data collection will be carried out, including when analyses have a defined hypothesis (*Route 1*), and when there is no defined hypothesis to test (*Route 2*). Secondly, we consider two scenarios where researchers plan to use existing data sources, including when analyses have a defined hypothesis (*Route 3*), and when there is no defined hypothesis to test (*Route 4*). For each scenario, we provide an example calculation.

Setting up a new study

When designing and setting up a study, researchers must decide on a wide range of study design parameters prior to collecting data. This includes decisions on including control groups, randomizing samples or participants, blinding researchers to conditions, adjusting for covariates, and collecting outcome measures within or between participants. However, the sample size that the study requires to answer the research question is often given relatively little attention. Often conventions within a field are used, rather than a comprehensive sample size calculation for each study (17, 18).

Researchers may want to test a defined hypothesis, such as whether giving patients a specific drug improves their health, compared to individuals not given the drug. Alternatively, the study may be an exploratory analysis or pilot study, where no hypothesis is specified, and no inferential statistics are calculated, but results must still be informative. For example, descriptive analyses may be required, or key parameters for use in future studies may be estimated. In both cases, sample size is critical to the precision of the effect sizes or other parameters estimated, but different approaches are required to determine the appropriate sample size.

Analyses with a defined hypothesis (*Route 1*):

Decide the minimum detectable effect size of importance

Once a hypothesis is defined, the first critical step is to determine the minimum detectable effect size of importance. This will depend on the research question and the subject area. An effect size may be deemed important because it is clinically, biologically or theoretically interesting. By implication, an effect size smaller than this is considered unimportant, or unlikely to change clinical practice. Sometimes, previous studies are used to inform the effect size of interest. However, caution should be taken where previous studies may be underpowered and the estimate may be biased (typically inflated) (8). This should be set prior to carrying out any formal power calculations.

It is also possible to calculate the minimum detectable effect sizes for a range of alpha levels (e.g., 5%, 1%) and power (e.g., 80%, 90%).

Carry out a sample size calculation

A sample size calculation relies on specifying the minimum power required, the alpha threshold chosen, the minimum detectable effect size of importance and an estimate of the variability of the outcome. Where the outcome is expressed in standardised units, such as a standard deviation, then the effect size and outcome variability will be combined in this standardised measure. Many resources are available online and in statistical packages, including Stata and R, to calculate the necessary sample size for a given effect size (19), with extensions for more complex models/estimates. Additionally, a number of articles and books provide detailed information on how to carry these out (20-22).

Adequate study power is conventionally defined as 80%, meaning that if a true effect of at least the size specified exists, there is an 80% chance the null hypothesis will be correctly rejected, with a 20% chance of a Type II error (i.e., incorrectly not rejected) (23). However, 80% is an arbitrary value and the optimal power for any given study will ultimately depend on the research question. For example, large scale clinical trials focused on confirming the efficacy of a new treatment may be designed with 90% or 95% power to detect a minimum detectable effect size of importance (9, 24, 25).

Assess whether the sample size is achievable

Researchers should then consider whether the sample size required is achievable. For example, are sufficient resources available, both financially and in terms of time to carry out the study as required? If the study requires recruitment of 1000 people with a given condition, how many such people are seen at the recruiting centre annually – and realistically, what proportion will be likely to agree to participate in the study? Further considerations and research options in this scenario are presented at the end of this article.

Example sample size calculation

In this hypothetical example, researchers hypothesised that normal weight individuals (BMI of 18-24.99 Kg/m²) have a different mean systolic blood pressure (mmHg) compared with overweight individuals (BMI of 25-30 Kg/m²). The literature indicates that a 10 mmHg lower systolic blood pressure reduces the risk of coronary heart disease by about 20% (26). In this example, using this risk difference in coronary heart disease as the motivation, the minimum detectable effect size of importance decided was determined to be 10 mmHg. The standard deviation in this example was set to 20, based on the standard deviation of systolic blood pressure recorded in UK Biobank participants (27). Note that this is a subjective decision, and other values, with different justifications, could have been chosen. In this example, the effect size is specified as a mean in each group (115 mmHg and 125 mmHg). However, it is the difference between the means that is important, and a mean of 5 and 15 would give the same sample size requirement for example. The sample size estimate is for a two-sample t-test for a difference in means, and a sample size calculation indicates that a total of 128 individuals (64 in each group) if recruiting equal numbers of normal weight and overweight individuals, would be required to detect a minimum difference of 10 mmHg (Table 1) with 80% power at a significance level of 5%. Example Stata code used to carry out this calculation, and all other examples presented throughout, are available in the supplementary material.

Table 1: Sample size required for a defined hypothesis to test for a 10mm Hg mean difference in systolic blood pressure between a group of normal weight and overweight individuals.

Study parameters	
Alpha	0.05
Power	80%
M1	115.00
M2	125.00
Delta (=M2-M1)	10.00
SD	20.00
Required sample size (N)	
N (if equal sized groups)	128
N per group (if equally sized)	64

Alpha = significance level; delta = minimum detectable effect size of importance; M1 = mean systolic blood pressure in group1; M2 = mean systolic blood pressure in group 2. SD = standard deviation of blood pressure, assumed the same in both groups. Hypothetical values based on clinically meaningful differences in blood pressure

The same sample size calculation could have been carried out using standardised measures of the outcome. Here, the minimum detectable effect size (10 mmHg) is half of the standard deviation (20 mmHg), and so the same calculation could use a standard deviation of 1 and an effect size of 0.5. However, when reporting sample size calculations, an effect size on a standardised scale is harder interpret than one on the natural scale of the outcome (here, mmHg).

Analyses without a defined hypothesis (*Route 2*):

Researchers may not have a defined hypothesis to test. For example, a study might be designed to estimate the BP difference between normal and overweight individuals, rather than test a null hypothesis that this difference is zero. Here, although a P-value for the null hypothesis may be calculated, the key analysis is to estimate the BP difference, with a confidence interval. The study size should be large enough such that (on average, if repeated across many studies), the confidence interval would be of a desired width.

Pilot and feasibility studies are other examples of study designs (28) where researchers may not intend to formally test hypotheses (29). By definition, P-values should not be calculated as there is no hypothesis to test. Any analysis of these studies should be largely descriptive or focus on estimating key parameters (with confidence intervals) that will inform the design of a future, adequately powered full study (30). These studies should be large enough to estimate relevant parameters with suitable precision to inform future research (31). The aim here is to have a large enough sample that (on average, if repeated across many studies) you would achieve the desired precision.

216 Determine the parameter of interest to estimate

217 For a feasibility study, the parameter to be estimated is defined in the study design (e.g., the mean

218 BP difference between normal weight and overweight individuals, in the first example above).

219 The first step when considering the size of a feasibility or pilot study should be determining the main

220 parameter of interest. As in all scenarios, this will depend on the research question being

221 investigated, but in these study designs it will likely depend on the parameters most important for

222 future studies, whether this is scaling up a pilot study or defining hypotheses. In a review of pilot and

223 feasibility studies, Lancaster and colleagues identified seven main reasons for carrying out a pilot

224 study, one of which was to obtain information that will guide the design of future studies. This

225 included collecting data to inform subsequent sample size calculations, and/or assess the integrity of

226 a study protocol and test randomization procedures (30). Data collected from the pilot study could

227 include the standard deviations of the outcome measure, or proportions in the control groups, to

228 inform future power calculations (30).

229 Decide the precision required and estimate the sample size required

230 The second step is to determine the precision with which this parameter should be estimated. For

231 example, a pilot study comparing two groups may be used to estimate the standard deviation of the

232 primary outcome measure, which can then be used to inform the design of an adequately powered

233 full study (31). This estimate of the standard deviation must be estimated sufficiently precisely to be

234 useful. Calculators are available which use a specified level of precision to calculate the required

235 sample size (13, 32), rather than estimate sample size based on statistical power and a hypothesis

236 test at a specified alpha.

237 Researchers should then consider whether the required sample size is achievable, and therefore

238 whether the research should be carried out or not. This is particularly important for a pilot or

239 feasibility study, when one of the aims may be to inform future power calculations (33). Results

should then be published according to EQUATOR guidelines, using relevant CONSORT reporting checklists for pilot and feasibility studies (34, 35).

Example sample size calculation

In this hypothetical example, a feasibility study was set up to determine likely participation into a full trial. The investigators hypothesised that 0.8 of eligible patients would agree to participate in the trial and wanted to estimate this with a precision of +/- 0.1 because of the impact of uncertainty on costings for the main trial.

Using the equation:

$$i) n = \frac{Z^2 x P(1-P)}{e^2}$$

where:

Z = the standard normal distribution value for the desired confidence interval

P = the true proportion

E = half the desired precision width.

The sample size calculation carried out indicates that a feasibility study size of 61 would be required to estimate participation with a precision of 0.1 (i.e., a confidence interval width of 0.2), if the true participation rate was 0.8 (Table 2).

Table 2: Sample size required to estimate participation proportion of 0.8 in a trial with a precision of 0.1

Study parameters	
Estimated proportion for participation	0.8
Alpha	5%
Precision (=1.96*SE)	0.1
Required sample size	
N	61

Alpha = significance level; SE = standard error

Using existing data

Secondary analyses of existing datasets are increasingly common, but here sample sizes are constrained by the size of the dataset. However, the researcher should still consider the effect size that can be detected (if testing a hypothesis) or the precision of an estimate that can be calculated, given the sample size. Although the sample size cannot usually be increased, understanding the effect sizes that can be detected with adequate power, or the precision with which key parameters can be estimated, will inform whether the study is likely to give a robust answer to the research question. As above, the researcher should be open to the possibility that the analysis should not be conducted if it is likely to be uninformative or should conduct the analysis in awareness of the limitations and seek additional datasets to increase power.

Analyses with a defined hypothesis (*Route 3*):

Decide the minimum detectable effect size of importance

The first step is to determine the minimum detectable effect size given the sample size available. As discussed for *Route 1*, whether this is clinically, biologically or theoretically interesting will depend on the research question and the hypothesis. Additionally, researchers should decide what level of power is adequate for the research question.

Identify the sample size available

When analysing secondary data, the variables required for the analysis may not be observed for all individuals, and samples sizes for analysis are often smaller than the size of the full dataset. This should be determined prior to any power calculations, so that the true sample size is known. For example, the Avon Longitudinal Study of Parents and Children is a birth cohort study in the UK that is extensively used for secondary data analysis (36). In total, 15,247 pregnancies were recruited, with genome-wide data collected on 8,365 children. Therefore, in any analysis using genetic data, the sample size would be just over half the original cohort. Additionally, not all children attend all follow-

up clinics, therefore any longitudinal analyses may have much smaller numbers (36). For an example of this, see Sallis and colleagues, where the analysis sample ranged from 3,289 to 5,480 individuals from the full sample of 15,427 individuals, resulting in analyses having up to 84% power to detect prespecified effect sizes (37). As per EQUATOR reporting guidelines, a flow diagram for eligibility and exclusions, starting at the full sample size and ending at the sample size for analysis should be used. Researchers should also consider whether methods for imputing missing data at this stage might be appropriate to maximise sample size (38).

Carry out a power calculation

A power calculation should be carried out to determine the power available to detect the previously determined minimum detectable effect size of importance, given the available sample size. Resources are available for this online, in statistical packages and elsewhere (19-22).

If a power calculation indicates that the minimum detectable effect size of importance is large, this may mean that the study has limited power to detect smaller but nevertheless clinically, biologically or theoretically interesting effect sizes.

Example power calculation

This example follows a similar hypothesis to the example sample size calculation in *Route 1*, where it was hypothesised that normal weight individuals have a different mean systolic blood pressure compared with overweight individuals. The minimum detectable effect size of importance was set at 10 mmHg and the minimum power required for meaningful inference to be 80%. In this example, the data has already been collected, and there is a total of 100 individuals (50 normal weight and 50 overweight individuals) available for analysis. Power is estimated for a two-sample t-test for a difference in means. This power calculation indicated that for a sample size of 100, this analysis would have 70% power (Table 3) to detect a difference of 10 mmHg at the 5% significance level. Using the same parameters, the minimum detectable effect size that can be detected with 80% power can be calculated, which in this case is a mean difference of 11.32 mmHg (Table 4). This value

is greater than the pre-specified minimum detectable effect size and consideration should be given as to whether these analyses should be carried out.

Table 3: Power to test for a 10 mmHg mean difference in systolic blood pressure between a group of normal weight (N=50) and overweight (N=50) individuals.

Study parameters	
N	100
N per group	50
Alpha	5%
M1	115.00
M2	125.00
Delta	10.00
SD	20.00
Power	
Estimated Power = 70%	

N = sample size; alpha = significance level; delta = minimum detectable effect size of importance; M1 = mean systolic blood pressure in group1; M2 = mean systolic blood pressure in group 2. Hypothetical values based on clinically meaningful differences in blood pressure SD = standard deviation of blood pressure, assumed equal in both groups

Table 4: Minimum detectable effect size for a mean difference in systolic blood pressure between a group of normal weight (N=50) and overweight (N=50) individuals with power of 80%.

Study parameters	
N	100
N per group	50
Alpha	5%
Power	80%
M1	115.00
SD	20.00
Estimated effect size and experimental group mean:	
Delta	11.32
M2	126.32

N = sample size; alpha = significance level; M1 = mean systolic blood pressure in group1; SD= standard deviation; delta = minimum detectable effect size of importance; M2 = mean systolic blood pressure in group 2. Hypothetical values based on clinically meaningful differences in blood pressure

Analyses without a defined hypothesis (*Route 4*):

Determine the parameter of interest

As previously discussed for analyses collecting primary data but not aiming to test a hypothesis (*Route 2*), researchers should initially identify which parameters need to be estimated, and then decide upon the desired precision with which they should be estimated

Identify the actual sample size available for analysis

In the case of secondary analyses of datasets, the sample size available for analysis may be much smaller than that of the whole dataset (see above). Researchers should determine the sample size prior to carrying out any further calculations.

Calculate the precision the available sample size will provide

As discussed previously for *Route 2*, it is important that the sample size can precisely estimate the parameters of interest. Precision calculators are available which allow researchers to determine the precision with which parameters of interest can be estimated from available sample sizes (32).

Assess whether the sample size can provide appropriate precision for the parameters of interest

As in all scenarios, researchers should again consider whether the sample size available will be able to estimate the parameters of interest with appropriate precision to be informative and decide whether the study can be carried out, or whether alternative approaches are necessary (e.g., identifying additional datasets that will allow the available sample size to be increased). Critically, for all scenarios, researchers should not go back to the parameters of interest at this stage and change the meaningful precision required in order to give the impression that the available sample size is adequate when it is not.

Example power calculation

This example follows a similar hypothetical example to *Route 2*. In this hypothetical feasibility study, researchers wanted to identify likely participation rates in future studies. Researchers used data from a previous study with a sample size of 80 participants, out of 100 invited (0.8 participation rate). Researchers wanted to estimate a confidence interval for the precision of this participation rate. They specified the minimum precision of interest to be 0.1.

Using the equation:

$$ii) e = Z \sqrt{\frac{p(1-p)}{n}}$$

where:

Z = the standard normal distribution value for the desired confidence interval (equivalent to 1.96 for a 95% confidence interval)

P = the true proportion

n = the fixed sample size

This calculation indicated for a fixed sample size of 80, where the participation rate (proportion of interest) was 0.8, the precision ($1.96 * SE$) is 0.09 (Table 5). This is smaller than 0.1, indicating this feasibility study meets the prespecified requirements.

Table 5: Precision available to test for a difference in risk of hypertension in normal weight compared with overweight individuals with a fixed sample size of 80

Study parameters	
Alpha	5%
Proportion	0.8
N	80
Precision of proportion rate	
Precision	0.09

Further considerations

Use of animals in research

Additional constraints on sample size may exist. Animal researchers are required to minimise the number of animals used; indeed, this is a legal requirement in the UK under the Animals (Scientific Procedures) Act 1986, and evidence of efforts to achieve this is often required by funders and research ethics committees. Whilst this is clearly important, the fundamental issue is that animals should not be sacrificed needlessly. Therefore, an underpowered study that cannot provide a robust answer to the research question is as unethical as one that is overpowered. The critical consideration is that the *right* number of animals is used to answer the research question, and that the research question is sufficiently important to justify the use of that number of animals. A series of underpowered studies, which may occur if none can provide a robust answer to the underlying research question, may use more animals than a single, adequately powered study (11). A number of resources are available on the National Centre for the Replacement, Refinement and Reduction of Animals in Research (NC3Rs) website, including the Experimental Design Assistant that supports power calculations (39).

Post hoc power calculations

It is widely accepted that *post hoc* (or observed) power calculations should not be carried out if based on the observed results, because observed power does not provide any additional information over the observed P-value. If calculating power *post hoc*, the sample size, effect size and the P-value are all constant values calculated from the study, and power becomes a direct monotonic function of the P-value (see Hoenig and Heisey (40) for a full explanation).

What to do when power cannot be achieved?

If the required sample size is deemed unachievable, or the available sample size does not provide adequate power, then a number of options are available, including not carrying out the study. In

order to increase the power for a given sample size and minimum effect size, more precise measures of the outcome could be used (e.g., taking the average of three repeated measures of blood pressure, rather than just one), or more powerful study designs used (e.g., within-subjects design where each individual receives both the intervention and control treatment – usually in random order, as in a cross-over trial – in a study) (41). Alternatively, it may be possible to identify additional sources of data that could be combined to increase the sample size (for example, consider a multi-centre study).

Another possibility is that the study could be carried out, but when reported the low power of the study made clear, with the caution that the results should be interpreted accordingly. This latter approach may be appropriate if little evidence on a particular hypothesis currently exists, and the results could potentially be included in a future meta-analysis. In this case it is critical to ensure that results are reported irrespective of the results (i.e., whether or not the hypothesis test is significant), in order to avoid contributing to publication bias. If an underpowered study does go ahead, the power of the study, or the minimum effect sizes the study had adequate power to detect should be reported. This allows readers to judge the robustness of the reported findings. The researchers should ensure that the results are made publicly available, regardless of their statistical significance. Ultimately, however, researchers need to be more open to the possibility that they simply may not be able to robustly answer a research question of interest with the resources available to them.

Conclusions

We have developed a flowchart (Figure 1) which outlines how sample size and power calculations should be carried out, considering different research scenarios. The sample size or power of a study should always be considered when designing the study, even if the sample size is fixed, or there is no defined hypothesis. When collecting new data, a sample size calculation should be carried out, specifying minimum effect sizes which could be detected. When using existing data, a power

- 416 calculation or precision calculation should be used. The rationale behind the sample size or power
- 417 calculation should be reported alongside the results.

References

1. Ioannidis JP. Why most published research findings are false. *PLoS Med.* 2005;2(8):e124.
2. Reality check on reproducibility. *Nature.* 2016;533(7604):437.
3. Lithgow GJ, Driscoll M, Phillips P. A long journey to reproducible results. *Nature.* 2017;548(7668):387-8.
4. Baker M. 1,500 scientists lift the lid on reproducibility. *Nature.* 2016;533(7604):452-4.
5. Hunter P. The reproducibility "crisis": Reaction to replication crisis should not stifle innovation. *EMBO Rep.* 2017;18(9):1493-6.
6. Maxwell SE, Kelley K, Rausch JR. Sample size planning for statistical power and accuracy in parameter estimation. *Annu Rev Psychol.* 2008;59:537-63.
7. Dickersin K. The Existence of Publication Bias and Risk-Factors for Its Occurrence. *Jama-J Am Med Assoc.* 1990;263(10):1385-9.
8. Button KS, Ioannidis JP, Mokrysz C, Nosek BA, Flint J, Robinson ES, et al. Power failure: why small sample size undermines the reliability of neuroscience. *Nat Rev Neurosci.* 2013;14(5):365-76.
9. Benjamin DJ, Berger JO, Johannesson M, Nosek BA, Wagenmakers EJ, Berk R, et al. Redefine statistical significance. *Nat Hum Behav.* 2018;2(1):6-10.
10. Thompson VA, Campbell JID. A power struggle: Between- vs. within-subjects designs in deductive reasoning research. *Psychologia.* 2004;47(4):277-96.
11. Button KS, Ioannidis JP, Mokrysz C, Nosek BA, Flint J, Robinson ES, et al. Confidence and precision increase with high statistical power. *Nat Rev Neurosci.* 2013;14(8):585-6.
12. Dumas-Mallet E, Button KS, Boraud T, Gonon F, Munafò MR. Low statistical power in biomedical science: a review of three human research domains. *R Soc Open Sci.* 2017;4(2):160254.
13. Rothman KJ, Greenland S. Planning Study Size Based on Precision Rather Than Power. *Epidemiology.* 2018;29(5):599-603.
14. Announcement: Reducing our irreproducibility. *Nature.* 2013.
15. Moher D, Hopewell S, Schulz KF, Montori V, Gotzsche PC, Devereaux PJ, et al. CONSORT 2010 explanation and elaboration: updated guidelines for reporting parallel group randomised trials. *BMJ.* 2010;340:c869.
16. von Elm E, Altman DG, Egger M, Pocock SJ, Gotzsche PC, Vandenbroucke JP, et al. Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) statement: guidelines for reporting observational studies. *BMJ.* 2007;335(7624):806-8.
17. Carter AR, Tilling K, Munafò MR. A systematic review of sample size and power in leading neuroscience journals. *bioRxiv.* 2017.
18. Lakens D. Sample Size Justification. *PsyArXiv.* 2021.
19. Kane S. Sample Size Calculator ClinCalc [updated 01/07/2017. Available from: <http://clincalc.com/stats/samplesize.aspx>.
20. Kirkwood BR, Sterne JAC, Blackwell S. Essential medical statistics. Malden, MA: Blackwell Science; 2017.
21. Noordzij M, Tripepi G, Dekker FW, Zoccali C, Tanck MW, Jager KJ. Sample size calculations: basic principles and common pitfalls. *Nephrol Dial Transplant.* 2010;25(5):1388-93.
22. Fox N. HA, and Mathers N. Sampling and sample size calculation. The NIHR RDS for the East Midlands / Yorkshire & the Humber 2007; 2007.
23. Batterham AM, Atkinson G. How big does my sample need to be? A primer on the murky world of sample size estimation. *Phys Ther Sport.* 2005;6(3):153-63.
24. Lakens D, Adolfs FG, Albers CJ, Anvari F, Apps MAJ, Argamon SE, et al. Justify your alpha. *Nature Human Behaviour.* 2018;2(3):168-71.
25. Collaborators WT. Effect of early tranexamic acid administration on mortality, hysterectomy, and other morbidities in women with post-partum haemorrhage (WOMAN): an international, randomised, double-blind, placebo-controlled trial. *Lancet.* 2017;389(10084):2105-16.

26. Lawes CM, Bennett DA, Lewington S, Rodgers A. Blood pressure and coronary heart disease: a review of the evidence. *Semin Vasc Med.* 2002;2(4):355-68.
27. Biobank U. 2019 [updated 07/09/2019. Summary statistics for systolic blood pressure automated reading in UK Biobank participants]. Available from: <http://biobank.ctsu.ox.ac.uk/crystal/field.cgi?id=4080>.
28. Porta M. A dictionary of epidemiology: Oxford university press; 2014.
29. Lee EC, Whitehead AL, Jacques RM, Julious SA. The statistical interpretation of pilot trials: should significance thresholds be reconsidered? *BMC Med Res Methodol.* 2014;14:41.
30. Lancaster GA, Dodd S, Williamson PR. Design and analysis of pilot studies: recommendations for good practice. *J Eval Clin Pract.* 2004;10(2):307-12.
31. Bell ML, Whitehead AL, Julious SA. Guidance for using pilot studies to inform the design of intervention trials with continuous outcomes. *Clin Epidemiol.* 2018;10:153-7.
32. Barrett M. precisely: An R package for estimating sample size based on precision 2018 [Available from: <https://malcolmbarrett.shinyapps.io/precisely/>].
33. Thabane L, Ma J, Chu R, Cheng J, Ismaila A, Rios LP, et al. A tutorial on pilot studies: the what, why and how. *BMC Med Res Methodol.* 2010;10:1.
34. Eldridge SM, Chan CL, Campbell MJ, Bond CM, Hopewell S, Thabane L, et al. CONSORT 2010 statement: extension to randomised pilot and feasibility trials. *BMJ.* 2016;355:i5239.
35. Thabane L, Hopewell S, Lancaster GA, Bond CM, Coleman CL, Campbell MJ, et al. Methods and processes for development of a CONSORT extension for reporting pilot randomized controlled trials. *Pilot Feasibility Stud.* 2016;2:25.
36. Boyd A, Golding J, Macleod J, Lawlor DA, Fraser A, Henderson J, et al. Cohort Profile: the 'children of the 90s'--the index offspring of the Avon Longitudinal Study of Parents and Children. *Int J Epidemiol.* 2013;42(1):111-27.
37. Sallis H, Evans J, Wootton R, Krapohl E, Oldehinkel AJ, Davey Smith G, et al. Genetics of depressive symptoms in adolescence. *BMC Psychiatry.* 2017;17(1):321.
38. Hughes RA, Heron J, Sterne JAC, Tilling K. Accounting for missing data in statistical analyses: multiple imputation is not always the answer. *Int J Epidemiol.* 2019;48(4):1294-304.
39. NC3Rs. Experimental Design [Available from: <https://www.nc3rs.org.uk/experimental-design>].
40. Hoenig JM, Heisey DM. The Abuse of Power. *The American Statistician.* 2001;55(1):19-24.
41. Mills EJ, Chan AW, Wu P, Vail A, Guyatt GH, Altman DG. Design, analysis, and presentation of crossover trials. *Trials.* 2009;10:27.

Figure 1: Steps required to determine the appropriate sample size or available power of a study

*For a continuous outcome additionally estimate the variability of the outcome