



This is a repository copy of *Why sample size estimates?*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/135070/>

Version: Accepted Version

Article:

Weber, E.J. and Hoo, Z. orcid.org/0000-0002-7067-3783 (2018) Why sample size estimates? *Emergency medicine journal*, 35 (12). pp. 755-756. ISSN 1472-0205

<https://doi.org/10.1136/emermed-2018-207763>

This article has been accepted for publication in *Emergency Medicine Journal*, 2018 following peer review, and the Version of Record can be accessed online at <http://dx.doi.org/10.1136/emermed-2018-207763>. © Authors (or their employer(s)). Reuse of this manuscript version (excluding any databases, tables, diagrams, photographs and other images or illustrative material included where another copyright owner is identified) is permitted strictly pursuant to the terms of the Creative Commons Attribution-Non Commercial 4.0 International (<https://creativecommons.org/licenses/by-nc/4.0/>)

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Emergency Medicine Journal

Why sample size estimates?

Journal:	<i>Emergency Medicine Journal</i>
Manuscript ID	emermed-2018-207763.R1
Article Type:	Concepts
Date Submitted by the Author:	30-Jun-2018
Complete List of Authors:	Weber, Ellen; University of California San Francisco, Emergency Medicine Hoo, Zhe Hui; University of Sheffield, School of Health and Related Research (SchARR); Sheffield Teaching Hospitals NHS Foundation Trust, Sheffield Adult CF Centre
Keywords:	statistics, research, methods, research, clinical, publication, epidemiology

SCHOLARONE™
Manuscripts

Review Only

WHY SAMPLE SIZE ESTIMATES?

Understanding Sample Size Calculations and Statistical Significance

EMJ asks that authors include in their Methods section a sample size calculation. Authors may wonder why we require this, and readers may wonder why they just shouldn't skip over it. In this paper, we explain what information can be gleaned from the sample size estimate, and along the way, the true meaning of statistical significance.

Why should I calculate a sample size?

In every experiment, you are using a sample to make inferences about a much larger target population. The problem is that you don't know if your sample is truly representative of the entire population. Your sample may be younger or older than the rest, or have different illnesses. Clearly, the larger your sample, the more of this variation you will capture, the more representative your sample is likely to be and the more precise your estimates. But you don't have infinite resources. So the goal of a sample size calculation is to help to ensure we have enough subjects to take account of that underlying variation in determining whether, for example, treatment a is better than treatment b and if so, by how much?

In a descriptive study, you estimate the sample size based on the predicted proportion of patients that have the illness or the outcome. In a comparative study, you are looking to see if there is a difference in some outcome between two or more groups. The groups may differ by one or more aspects (e.g. one group is < 65 years and another is ≥ 65 years; or one group had training and another didn't, or one group got the drug and another didn't, as in a randomised trial.). The goal of a sample size calculation here is get a fairly precise estimate of the effect in each group, with reasonable assurance that you didn't get these results by chance. This is where we get into statistical significance.

Statistical Significance – What does it really mean?

Many people misinterpret what “statistically significant” means. Don't feel bad if you're one of them – the concept comes from the language of hypothesis testing, which is not something that is very intuitive to clinicians. Using the language of hypothesis testing, statistical significance means that the probability of rejecting the null hypothesis (that there is no effect of the intervention) when it is true is low. Translation: **Statistical significance means that result you got is unlikely to be due to chance. Importantly**, it does *not* mean that the numbers for each group are greatly different from each other, or that the difference is meaningful clinically.

Remember, you began with a sample of the target population. To be more certain of your result, you'd want to repeat the study on different samples in that same population. Statistical significance means that if you repeated the study multiple times, with different samples from the target population, you would get very similar results most of the time. Your answer is unlikely to be a fluke!

1
2
3 . Statistical significance is typically inferred from p values. To interpret a p-value, it is helpful
4 to remember that the p value comes from the world of hypothesis testing, where the
5 experiment is set out such that the goal is to determine if you can reject the null hypothesis.
6 In a randomized study comparing two treatments for example, the null hypothesis is that
7 there is no difference between the treatments. If you find a difference, then the p value tells
8 you whether you have sufficient statistical evidence to reject the null hypothesis. A smaller p
9 value implies a greater statistical incompatibility with the null hypothesis¹ In other words, the
10 smaller the p value, the stronger the evidence for rejecting the null hypothesis.
11

12
13 Let's say you did an experiment to compare two variables and you obtained a set of
14 values from that experiment. You then compare those variables using a statistical
15 test, which generates a p-value. Let say that p-value is 0.03. The means the
16 probability of obtaining that value (or one more extreme i.e. more different) by
17 chance in that single experiment is 0.03. If you have specified a "significant" p-value
18 as 0.05, it means the p-value you obtained is "statistically significant", i.e. unlikely
19 you have obtained that set of results just by chance. But if you have specified a
20 "significant" p-value as 0.01, that means the p-value you obtained is "not statistically
21 significant".
22
23

24
25 In most studies, statistical significance (alpha) is set at .05. This means that the
26 probably of getting the result you have gotten, or one more extreme, by chance is
27 5%. If you would like to be even more certain the result isn't due to chance, you set
28 alpha at .01. Note that *this doesn't mean the magnitude of the effect is bigger*, only
29 that you are more certain. In reality, when you do an experiment you calculate a p
30 value, which may be greater or less than 0.05 or 0.01. If you have set your alpha at
31 .05, and your p value is, say .03, you can call your results statistically significant; that
32 is, the probability of getting the result you have gotten, or one more extreme, is
33 within the range of chance you're willing to take.
34
35

36 *Getting back to Sample Size –*

37 Sample size for a comparative study is determined by the size of the effect you are
38 looking for, and variation in the population and two other parameters – the level of
39 statistical significance (alpha) and the power (1-B) . We've already said that by
40 convention, 0 .05 is usually chosen as the cut-off for statistical significance. Power is
41 the likelihood that we won't reject the null hypothesis when it is actually false.
42 Translation: Power is the likelihood of finding an effect when there is one. The
43 higher the power (and lower the Beta), the greater the likelihood. When estimating
44 the effect size, we may use parameters from other studies to guide us or we may
45 choose an effect size we wish to see (below which, say, the cost of a treatment isn't
46 worth it). The variation generally comes from other studies. These are of course
47 estimates, and you don't know if the variation in your sample will be the same, and
48 that can mean even if you reach the estimated sample size, and see an effect, it might
49 not reach statistical significance.
50
51
52

53
54 It is always preferable to base your sample size on the desired effect size.
55 Sometimes this may result in an unrealistically large sample size. In this case,
56 researchers may adjust the significance level, or the power. It is tempting, but
57
58
59
60

1
2
3 probably a bad idea, to adjust the effect size. It might not seem obvious at first, but
4 looking for larger treatment effects generally means needing smaller sample sizes: if
5 the effect of a treatment is very large, you are likely to see it very early in your
6 experiment, and this would suggest that you don't need a very large sample. It may
7 be quite tempting to design a study this way to shorten recruitment time and lower
8 your costs, but few things (in real life) have large effect sizes, so if you decide on too
9 large an effect size, even if a clinically important differences exists, a "statistically
10 significant" difference could not be detected with that sample size. Moreover,
11 you've lost some researcher integrity here in no longer sticking with your
12 hypothesized effect size.
13
14
15

16
17 Failing to show a significant difference when there might be one is called (in stats
18 speak) a Type II error. Ultimately it boils down to a "too small" sample size, but this
19 could be the result of a smaller effect or more variation than anticipated. This can
20 happen in any experiment – a sample size estimate is just that, an estimate. But it's
21 the best tool we have.
22
23

24
25 This leads to the corollary, which is that with a large enough sample size, even very
26 small differences that are of no clinical importance can be detected because the
27 likelihood of getting a result by chance decreases with more subjects. This is typical
28 of large database studies with thousands of patients. Here the integrity and
29 knowledge of the researcher and the scepticism of the reader is tested in
30 determining if the statistically significant finding is clinically meaningful. So while
31 we care about statistical significance from the point of view of how certain we are
32 that the results did not occur by chance alone, we must keep our eye on the target –
33 which is whether the difference is clinically meaningful.
34
35
36

37 A researcher who performs a study without a sample size estimate but finds a p
38 value $< .05$ may say: "Well, I got a significant p value, so that must mean I had a large
39 enough sample size" The problem with this is, just as we discussed above, if a study
40 is under-powered, it will only be able to detect a large effect. Should an under-
41 powered study happen to find a significant difference, it is likely that effect is
42 inflated (i.e. the difference found in the study is higher than the actual difference)
43 and the results are less likely to be reproducible.^{2,3}
44
45
46

47
48 Although we often think of statistical significance as a black and white issue, and
49 tend to completely ignore any results that are not significant, there is much more to
50 interpreting results than that. Moreover, many statisticians will argue that the
51 sample size calculation should not be put in the Methods section of the final paper,
52 as there are other ways for the reader to determine if the author has reached the
53 required sample size. This leads us to a later instalment on confidence intervals.
54
55
56
57
58
59
60

References

1. Wasserstein RL, Lazar NA. The ASA's statement on p-values: context, process, and purpose. *Am Stat* 2016;70:129-33.
- 2 Button KS, Ioannidis JP, Mokrysz C, et al. Power failure: why small sample size undermines the reliability of neuroscience. *Nat Rev Neurosci* 2013;14:365-76.
3. Ioannidis JP. Why most discovered true associations are inflated. *Epidemiology* 2008;19:640-8

Acknowledgments:

The authors would like to acknowledge Dr. Jason Oke for his helpful input.