Study Planning: Sample Size

Project team

Gjalt-Jorn Ygram Peters
Peter Verboon

Open Educational Resource van de Open Universiteit

This document is in progress. It is intended as guideline for use by ethical review committees and institutional review boards, researchers, and students. We may turn it into a paper at some point. For now, we just make it accessible through PsyArXiv.

This document is released under the [CC-BY-NC-ND 4.0 International licence](http://creativecommons.org/licenses/by-nc-nd/4.0/) (URL: <http://creativecommons.org/licenses/by-nc-nd/4.0/>). It can be distributed freely, but cannot be adapted ("remixed") or sold. This version of the document has been generated on 2017-05-30 at 17:14:26. When citing this document, use this reference:

Peters, G.-J.Y. & Verboon, P. (2017). *Study Planning: Sample Size*. Open Universiteit, Heerlen, The Netherlands.
Downloaded from: <https://osf.io/preprints/psyarxiv/94m27> (URL: <https://osf.io/preprints/psyarxiv/94m27>)

This document is primarily intended as guideline for use by students and researchers who are in the process of planning their research and need to establish the sample size they require. Secondly, ethical review committees and institutional review boards can use this document as guidelines when determining whether a proposed study has any chance of yielding useful results (after all, studies with too few datapoints cannot teach us anything, regardless of whether the results are as the researchers expected or not). Thus, sample size planning is one of the most important decisions in the study planning process. End up with too low a number, and your study has failed even before it began (with too few participants, it is very likely that you find nothing where associations exist, or find associations where none exist). End up with too high a number, and precious resources are wasted (your time, participants' time, often money, and maybe most importantly: participants are scarce, and your study will use up participants that could otherwise have participated in other research). This document is intended to facilitate meaningful study sample size planning by offering relatively straightforward guidelines.

Background (refresher)

Populations

The goal of scientific research is to learn about reality. If people are studied, the goal is to learn about a population. This population is defined as all people that meet relevant criteria, for example, all high school students in the Netherlands, all migrant sex workers in Amsterdam, or all patients with a chronic disease who exercise insufficiently. The population also has a definition *over time*: to prevent data from being obsolete the moment data collection finished, populations are assumed to stay stable for a given period of time. The length of this period depends on the characteristics of the population: the population of 'first year psychology students' has a lower expiry date than the population of 'grandparents'.

Populations are almost always defined as having infinite size. This is sensible because they are usually so large so as to be practically of infinite size, and because they are not defined as all people satisfying the criteria at any given moment, but rather over a prolonged time period, the exact number cannot be established. If you do a study with a population that is instead finite (i.e. *not* defined as infinitely large), this document is currently not useful for you. In that case, contact a statistician.

Samples

Because populations are defined as infinitely large, it is impossible to make any statements about them with certainty. Data collection would, by definition, take infinitely. The only method available to make statements about the populations (and therefore, about reality), is by collecting data in a finite group of people: a sample. There are many ways to obtain such a sample from a population (or: to 'sample'). These differ in one crucial respect: representativeness.

Random samples

If the goal of the study is to obtain results that are representative of the population, this requires random selection of the sample. Random selection means that either every individual in the population has the exact same likelihood of being sampled, or that the likelihood differs for (categories of) individuals but is known, enabling weighing the cases in the sample. Of these two, the latter is relatively rare: in most random samples, participants from the population are simply randomly selected.

Quantitative research almost always requires random sampling for the underlying statistical methods to apply.

Nonrandom samples

If the goal is not to obtain results that are representative of the population, but to obtain, for example, an exhaustive list of characteristics, random sampling is very inefficient. In such cases, it is more efficient to select population members to maximize heterogeneity: for example, striving for different age groups, cultural backgrounds, and gender identities.

For qualitative research, as well as cognitive interviewing, nonrandom sampling is almost always the better choice.

Attrition and optimism

One important thing to keep in mind when planning studies is that whatever estimate is planned for, this will almost invariably be an underestimate of the sample size required for the study to succeed. There are a number of reasons for this which will be listed briefly. First, participants often drop out of studies. Second, participants

sometimes provide data that is not useable (e.g. unrealistic values), in which case they have to be excluded for some or all analyses. Third, participants may exhibit more variation (i.e. be more different) than expected, which inflates the error and effective sample size. Fourth, mistakes can be made which cause loss of data. Therefore, it is important to multiple the results of the planning with a given proportion. If no other guidelines are available, adding 20% seems reasonable.

Types of studies

The type of research question that a researcher wants to answer, or the hypothesis that they want to test, determines the method they use. In this document, four basic methods are distinguished. Because sample size planning mostly concerns quantitative research, that method will receive much more attention.

One thing is important to remember. Any single study is fallible. This fallibility decreases as sample size increases and as the quality of planning increases. Therefore, strong conclusions require either synthesis of many separate studies or very, very large samples. The latter is often unattainable in practice, and the first requires a relatively mature literature (that is preferably free from publication bias). Nonetheless, in small studies of a few hundred participants, strong conclusions cannot be justified, and sometimes even deemed unethical. Therefore, researchers desiring to draw conclusions sufficiently strong to, say, warrant a press release, should commit to collecting evidence affording such strong conclusions, and therefore, either conduct a systematic review and meta-analysis, or collect a sample that is much larger than what is required if weak conclusions suffice.

Systematic reviews and meta-analysis

The strongest conclusions can be drawn on the basis of syntheses of many separate studies. If the studies are quantitative in nature and all provide data enabling computation of effect size estimates for the same association, a meta-analysis can be conducted. Otherwise, researchers will have to settle for a systematic review. However, in systematic reviews and meta-analyses, researchers almost always collect *all* existing studies. No sample size calculations are therefore necessary; and because no participants are recruited, not ethical approval is necessary.

Pilot studies

When researchers want to conduct empirical research (sometimes known as 'primary research', to distinguish it from studies using existing documents as data such as systematic reviews and meta-analyses, which are then referred to as 'secondary research'), it is often desirable to first conduct a small scale study to test the feasibility of the intended study. Such small scale preliminary studies are known as pilot studies. There exist two types.

Pilot studies to establish feasibility of the procedure

When conducting a pilot study to ensure that a planned procedure is feasible, the required sample size depends on the number of potential problems that one needs to detect and how likely each problem is to occur. There is an excellent introduction available in Viechbauer, Smits, Kotz, Bude, Spigt & Crutzen (2015). This paper also introduces a free tool, which is available at pilotsamplesize.com (URL: pilotsamplesize.com) .

As a guideline, If a problem occurs with 5% probability in each potential study participant, the problem will almost certainly be identified (with 95% confidence) in a pilot study including 59 participants.

Pilot studies to obtain an effect size estimate to properly power a study

Sometimes, researchers want to conduct a pilot study to have an idea as to the population effect size so that they can perform power analyses for their main study with more certainty.

Such pilot studies typically require very large sample sizes, because with low sample sizes, the obtained effect size estimate is very unreliable. Because the sampling distribution from which effect size estimates are drawn are very wide for low sample sizes, the sample value for an effect size is unlikely to be close to the population value.

The objective, therefore, is to obtain sufficiently narrow confidence intervals to have some kind of certainty as to

likely population values for the effect size.

Tables with guidelines for simple scenario's are available in two open access publications (Moinester & Gottfried, 2014, for correlations, and Peters & Crutzen, 2017, for estimates of Cohen's d values, e.g. for experiments).

For more complicated (i.e. multivariate) designs, simulations will be necessary, because in multivariate designs, sampling distributions are conditional upon all variables in the model, so in these cases, no guidelines can be provided (although the guidelines for Cohen's d are probably useful to follow).

As can be seen in Table 1 of Peters & Crutzen (2017), assuming 95% confidence intervals are used, if a pilot has a sample size of 100, and a Cohen's d value of 0.5 is observed (a moderate effect size), the confidence interval includes all values from 0.12 to 0.91. Assuming that the point estimate or upper bound of the confidence interval represent the population value is hard to justify. Therefore, one would normally power the main study for a Cohen's d value equal to the lower bound, in this case 0.12. For 80% power, the power analysis would indicate that 2354 participants would be needed in total.

If instead a Cohen's d of 0.4 was observed in a pilot study with a sample size of 100, the confidence interval includes all values from 0.02 to 0.8, and the main study would have to be powered at 128148, which will often be impossible.

If a pilot study has a larger sample size, of course tighter confidence intervals are obtained. For example, if a pilot study has a sample size of 200 participants, and a Cohen's d value of 0.4 is observed, the confidence interval would include all values from 0.13 to 0.68, and detecting an effect size of the main study would have to be powered at 1998.

As will be clear from these examples, if a pilot study is used to determine the required sample size to power the main study, the pilot study will need to yield accurate estimates of that effect size, which means the pilot study requires a large sample size itself. This is because the sampling distribution from which the Cohen's d values obtained in a study are drawn are very wide:

Figure 1. The sampling distributions of Cohen's d for population values of 0.5 and 0.4 and sample sizes of 100 and 200 participants. The confidence intervals that would be obtained for sample point estimates that happen to correspond exactly to the population value are shown in blue.

For this reason, pilot studies are often not useful for determining the effect size to expect. Pilot studies are generally used for their first purpose: to establish whether a designed procedure is feasible, and to detect unforeseen problems. In those cases, the sample size computations discussed above can be used (e.g. 59 participants if the goal is to identify a problem expected to occur in at most 5% of the participants).

Cognitive interviews

Because cognitive interviews are not part of the curriculum (yet), they will likely be unfamiliar to many students, and therefore, will be explained a bit more in-depth.

Cognitive interviews enable researchers to establish the so-called cognitive validity of an operationalisation. An operationalisation has high cognitive validity if all presented stimuli are interpreted by the participants as intended by the researcher, and if the registration of responses is consistent with the intended responses of participants. This abstract definition is easier to understand when applied to examples of operationalisations, and of those, questionnaires are the easiest example. Questionnaires contain at least two types of stimuli: prompts to trigger psychological processing in the participants (questions) and labels corresponding to answer options, that should together cover the potential responses of participants to the prompts. A very simple example is an operationalisation of sex:

What is your sex?

- ☐ Female
- ☐ Male
- ☐ Other

Here, the first stimulus is “What is your sex?”. This is combined with three response options that are designed to prompt a behavior. These response options are each paired with an additional stimulus (the labels of the three response or answer options). This operationalisation is designed to trigger a process where the participants determine their sex, and then indicate their perceived sex (which, in this case, will almost always be the same as their biological sex) using one of the three response options.

For participants who speak English and who are familiar with questionnaires, the cognitive validity of this question is likely to be very high. For participant who do not speak English or who have not yet learned how they are supposed to respond, the cognitive validity will be lower.

Another example comes from the Hospital Anxiety and Depression Scale (HADS, where one of the questions is:

I get a sort of frightened feeling like ‘butterflies’ **in** the stomach.

- ☐ 0 Not at all
- ☐ 1 Occasionally
- ☐ 2 Quite Often
- ☐ 3 Very Often

This item has high cognitive validity in some cultures and almost no cognitive validity in others (see e.g. Maters, Sanderman, Kim & Coyne, 2013, for a discussion).

Cognitive validity applies not only to questionnaires, but also to manipulations: all operationalisations have a

degree of cognitive validity, and this cognitive validity is not a property of the operationalisation itself, but can differ in different populations, and in different samples. If the stimuli or response registration are interpreted differently by the participants than intended by the researcher, this means that the operationalisation is unlikely to be valid. Cognitive validity, therefore, is a first condition for validity.

This means that it is wise to establish cognitive validity of the intended operationalisations before starting large-scale data collection. Cognitive interviews are an efficient means of doing so.

There are two common situations where researchers will want to ensure the cognitive validity of an operationalisation:

- When the researchers use existing (potentially already validated) operationalisations in a population where cognitive validity of the operationalisations has not yet been established;
- When the researchers are developing a new operationalisation.

In both cases, the procedure is the same, although less participants will be needed for instruments that already have cognitive validity, which may be the case for some validated operationalisations.

For cognitive interviews, the general recommendation is to use 6-12 participants (Willis, 2005). However, cognitive interviewing is an iterative process, so multiple rounds may be necessary (and more rounds should be expected for newly developed operationalisations than for operationalisations that have already been thoroughly studied). In addition, it is wise to use a stopping rule, for example based on the saturation criterion of Francis, Johnston, Robertson, Glidewell, Entwistle, Eccles & Grimshaw (2010; and for more information on cognitive interviews, see e.g. Drennan, 2003; Haeger, Lambert, Kinzie & Gieser, 2012; Willis & Artino, 2013; or Willis, 2005).

It is important to note that cognitive interviews, although they use qualitative methods, are not qualitative research. There are two important differences. First, in cognitive interviews, the participants are not the object of study: the operationalisation under study is. There is no need to collect any data about the participants. Instead, participants are consulted as 'experts' on their own cognitive processing of the operationalisation, similar to how many Delphi studies involve participants. This means that participants' privacy is practically secured, unless researchers do decide to collect data about the participants themselves. Second, because no typical qualitative analyses need to be conducted, cognitive interviews usually do not need to be recorded. Instead, notes often suffice. This means that the raw data, being intrinsically anonymous, can easily be made public.

Qualitative versus quantitative research

Although quantitative and qualitative research were traditionally seen as 'epistemologically irreconcilable', it has been argued extensively that where it concerns psychology and education science, this reconciliation is possible and in fact desirable (e.g. Leppink, 2017; Shaw & Flowers, 2012; Hiles, 2012). Both methods can answer different types of research questions.

Quantitative research uses operationalisations to quantify aspects of reality. This quantification makes it possible to model the measurement and sampling errors, which make it possible to establish how likely it is that patterns observed in a given sample are due to chance, or represent patterns in the population. This procedure is both a strength and weakness of quantitative methods.

It is a strength because without modeling the sampling and measurement error all patterns observed in a sample may be purely due to chance. Being able to determine how likely it is that patterns in a sample (e.g. associations) are caused by chance enables relatively confident statements about the population (i.e. the reality) based on samples alone.

It is a weakness because this quantification requires well-calibrated operationalisations. Such operationalisations are designed to influence (if the operationalisation is a manipulation) or measure (if the operationalisation is a measurement instrument) strongly bounded aspects of reality, which of course have to be identified before data collection starts. In other words: quantitative research can only study patterns in variables that were already known or at least hypothesized in advance. This is a substantial limitation: if insufficient is known about a given variable, quantitative research is simply impossible because no valid and reliable operationalisations will exist. These will then first have to be developed.

Qualitative research does not require quantification, and in fact, pure qualitative research is inconsistent with quantification. There are two reasons for this.

First, quantification by definition requires valid and reliable operationalisations. If those are available, quantitative methods are preferable over qualitative methods, because quantitative methods allow the researcher to draw much stronger conclusions about the population due to their ability to model the sampling and measurement errors.

Second, in qualitative research, very rich data are collected, and these data are analysed in a very thorough detailed fashion to enable in-depth processing of all aspects of these rich data. The data collection procedures are therefore almost invariably very labour-intensive, as are the iterative analysis procedures. This means that on the one hand, given the amount of data collected from each participant, relatively few participants are typically needed; but simultaneously, collecting data with many participants is often not feasible.

Although this works out nicely (in qualitative research, the data collection and analysis procedures preclude collecting data from many participants, but the richness of the obtained data often eliminates the need for large samples), these small sample sizes mean that although sampling and measurement error cannot be modeled, they will be very, very large. That means that any pattern observed in the sample may very likely be due to chance alone, and the researchers can never assume that patterns in a sample are indicative of similar patterns existing in the population. If the relevant confidence intervals are reported as well, this practice becomes more defensible; but even then, the width of the sampling distributions from which the relevant statistics are drawn mean that even these confidence intervals can vary considerably from sample to sample.

In addition, in qualitative research, samples are often not randomly selected from the population: instead, participants are often deliberately selected, which means that the sample is not representative of the population, and therefore, statistics computed from the sample do not generalize to the population.

Therefore, in qualitative research, reporting numbers can rightfully be considered a questionable research practice.

This means that any patterns observed in qualitative research may not tell the researcher much about the population under study. However, at the same time, qualitative research is the only means available to a researcher studying a new area.

Qualitative research, therefore, is useful when a researcher desires to explore a new area, for example, to explore potential aspects of a psychological construct, or reasons people have for behaving in a certain way.

Quantitative research is useful when a researcher desires to study whether two or more variables are associated to each other in a population, how often something occurs in the population, or what the distribution is of some construct in a given population.

In addition, if a researcher desires to test a hypothesis, quantitative methods are always required, because the

conditions under which the collected data are considered to confirm or refute the hypothesis will always require statements about the population, which in turn require knowing the relative magnitude of patterns observed in a sample compared to the sampling and measurement error. In addition, such conditions are almost always formulated in quantitative terms, and therefore, require operationalisation of the relevant variables.

It is important to realise that the research questions a researcher wishes to answer determine the method that has to be used. It is not possible to study any research question with any method. Most students, for example, desire to learn about associations between variables. This always requires quantitative research and statistical analyses.

Qualitative research

Because qualitative research does not use quantification, there exist no methods for computing the sample size in advance. This means that the subjective nature of the study sample size planning process is much more salient for qualitative research.

Individual interviews

A compilation of expert voices and early career reflections yielded a number of determinants of required sample size, such as the breadth of the research questions, whether the focus is on identifying shared characteristics or differences in characteristics, and the time available (Baker & Edwards, 2012). Mason (2010) provided a useful overview and analysis of often used sample sizes, finding mean and median values of 31 and 28, respectively, with a very high standard deviation of 19, consistent with the general consensus that 'it depends'.

In qualitative research, much more than in quantitative research, it is common to adjust the sample size during data collection. In fact, this is considered good practice: the researcher wishes to achieve 'data saturation' which is the name of the situation where sampling additional participants does not yield additional insights. This is a vague concept, which thankfully has been operationalised by Francis, Johnston, Robertson, Glidewell, Entwistle, Eccles and Grimshaw (2010).

Following these recommendations, the following guidelines can be used. For clearly bounded research questions that concern relatively topics with limited scope, it seems wise to count on requiring a sample of at least 10 participants. However, most researchers study research questions that have a moderate scope, involving more than one psychological construct, and potentially concerning multiple relevant scenarios. In such cases, 20 or 30 participants may be a better initial sample size.

After the initial sample size has been reached. An operationalisation of data saturation can be employed to continue data collection until data saturation has been achieved. For example, researchers can keep sampling in sets of three participants, and stop data collection as soon as such a set yields not additional insights.

Note that although some mechanism for establishing data saturation is almost always employed, this does not mean that study sample size planning can usefully set the anticipated required sample size at a low value.

Although on the one hand, such a procedure is protected against insufficient data collection by a sensible data saturation principle, on the other hand, study sample size planning enabled the researcher to prudently plan the entire study. Reserving too little time for data collection can result in the undesirable scenario that data collection has to end before data saturation has been reached.

Focus groups

For some research questions, individual interview may not be the best approach, and instead, focus groups are used. This is the case, for example, when the topic under study is such that observing social interactions provides

useful data. Sometimes, focus groups are used if research questions are very simple as a means to collect data more efficiently. For focus groups, an additional consideration is that the dynamics of each group of participants will in part determine the data collected in that group. In addition, collecting in-depth data for all participants will be impossible, and therefore, focus groups are most adequate for studying psychologically superficial topics.

These additional dynamics of focus groups mean that researchers need to establish in advance how many focus groups they expect they will require. This number can be informed, for example, by the homogeneity of each focus group (more heterogeneous group may exhibit more between-group variation) and factors such as the degree of stigmatisation expected to surround the research question.

Quantitative research

Null Hypothesis Significance Testing (NHST)

The most familiar approach to quantitative research is ironically also the most debated approach, known as Null Hypothesis Significance Testing (NHST). In this approach a null hypothesis, almost always that two or more variables are *not* associated to each other, is used to construct a sampling distribution of an appropriate statistic. This sampling distribution is then used to compute the probability that the sample point estimate for that statistics is obtained in a given study. This probability is known as the p-value. Researchers then compare this p-value to a predetermined value (known as 'alpha', and almost always set at .05), and if the p-value is lower, this is interpreted as evidence that it is unlikely that the sampling distribution that was used to compute the p-value is in fact the sampling distribution from which the relevant statistic was drawn. Since this sampling distribution was constructed based on the null hypothesis, this results in rejection of the null hypothesis, which is often taken as evidence that the studies variables are associated in the population. Such a result is referred to as a 'significant' effect.

Much has been said about the problems of this approach (see e.g. Cumming, 2014, for an accessible introduction), but it is the only approach most researchers are familiar with, and is thus widely applied.

Free tools for power analysis

The analyses to determine the required sample size in an NHST framework allow researchers to compute the sample size necessary to obtain the desired power. Power is defined as the likelihood of obtaining a statistic corresponding to a p-value of less than the chosen alpha value (usually .05) assuming that the population effect size has a specified magnitude. For example, it is possible to compute the sample size required to achieve a power of 95% if in the population, two variables are associated with a Cohen's d of 0.5 (a magnitude tentatively labelled as 'moderate' or 'medium' by Cohen, 1988, 1992). If in the population, the effect size does indeed have that specified magnitude, a study of the sample size obtained in such so-called power analyses has 95% likelihood of obtaining an effect size estimate that occurs in less than alpha percent (e.g. less than 5%) of the studies given the null hypothesis sampling distribution. For such effect size estimates, the p-value will be lower than alpha (e.g. $< .05$), which the researcher will deem significant, and which will lead to rejection of the null hypothesis.

For simple designs, free tools are available such as G*Power (see <http://www.gpower.hhu.de/en.html> (URL: <http://www.gpower.hhu.de/en.html>)) and the pwr package for R. However, these tools only allow power analyses for bivariate designs, or, when power analyses are conducted for multivariate designs, the analysis often implicitly assume that the independent variables are not correlated. That is usually an untenable assumption. Therefore, more complicated designs will often require simulations. In some cases, though, adjustments for the correlations between predictors is possible, and simulations are not needed.

Collinearity in multivariate analyses

If researchers conduct multivariate analyses (here defined as analysis involving more than two variables simultaneously), the predictors will almost always be correlated. Because those independent variables are associated, and therefore can be predicted from each other, less data are available than if they had been independent. Therefore, the effective sample size is lower.

On [his website](http://davidakenny.net/cm/mr.htm#Multi) (URL: <http://davidakenny.net/cm/mr.htm#Multi>), Kenny explains that in mediation analyses, the effective sample size that remains can be computed:

$$N_{\text{effective}} = N_{\text{total}} (1 - r^2)$$

Where r^2 is the correlation between two predictors. This formula can be inverted to obtain a multiplier that can be used to correct the outcomes from power analyses based on independent predictors:

$$N_{\text{required}} = \frac{N_{\text{power analysis}}}{1 - r^2}$$

In multivariate analyses with more than two predictors, an estimate of this explained variance between predictors can be obtained, for example, by regressing each predictor on the others, converting the obtained R^2 values to regular multiple correlation coefficients by taking the square root, converting these to Z -values using Fisher's Z transformation, averaging them, and then converting the average back into a correlation coefficient, which is then finally squared again to obtain the average proportion of explained variance (R_{average}^2).

Multiple testing

Inherent to the nature of the NHST procedure, the chosen value of alpha also determines how often a p-value will erroneously lead to rejection of the null hypothesis. After all, by definition, in alpha percent of the studies a value of a statistic is obtained that occurs in alpha percent of the sampling distributions. For example, when researchers set their alpha at .05, if the null hypothesis is in fact true, 5% of the studies will obtain a value for the statistic that causes the null hypothesis to be rejected nonetheless. This is known as a Type 1 error.

Every p-value that is computed increase the likelihood of a Type 1 error. This is true regardless of whether those p-values are computed in multivariate or bivariate analyses (Cramer, van Ravenzwaaij, Matzke, Steingroever, Wetzels, Grasman, Waldorp, Lourens & Wagenmakers, 2014).

This means that researchers who choose to use NHST need to conduct their power analyses using an alpha of less than their desired probability of making a Type 1 error. For example, if a researcher will compute twelve p-values in their analyses, in their power analyses, they need to specify an alpha of for example $.05 / 12 = 0.004$. Other, more liberal, methods to correct this alkha value are available; but given the fact that researchers need to plan their sample size on the conservative rather than the liberal side (also see section "Attrition and optimism" above), it seems prudent to stick to this simple correction.

More complicated designs

Computing the required sample size for an NHST study is relatively straightforward for simple designs. There are two important considerations that are outlined below. For more complicated designs it is necessary to write a simulation, which will usually require involving a statistician in the study planning process.

Parameter Estimation

If the goal of a study is to determine an effect size, the objective is to obtain a confidence interval that is sufficiently narrow. Because the sampling distributions of effect sizes are usually known, these computations are relatively

straightforward using an iterative procedure.

Tables with guidelines for simple scenario's are available in two open access publications (Moinester & Gottfried, 2014, for correlations, and Peters & Crutzen, 2017, for estimates of Cohen's d values, e.g. for experiments).

For more complicated (i.e. multivariate) designs, simulations will be necessary, because in multivariate designs, sampling distributions are conditional upon all variables in the model, so in these cases, no guidelines can be provided (although the guidelines for Cohen's d are probably useful to follow).

Bayesian Inference

This section will be added later.

References

- Baker, S. E., & Edwards, R. (2012). How many qualitative interviews is enough ? National Centre for Research Methods Review Paper, 1–42. <http://doi.org/10.1177/1525822X05279903> (URL: <http://doi.org/10.1177/1525822X05279903>)
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ, NJ: Lawrence Erlbaum Associates.
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112(1), 155–159.
- Cramer, A. O. J., van Ravenzwaaij, D., Matzke, D., Steingroever, H., Wetzels, R., Grasman, R. P. P. P., ... Wagenmakers, E.-J. (2016). Hidden Multiplicity in Multiway ANOVA: Prevalence, Consequences, and Remedies. *Psychonomic Bulletin & Review*, 23, 640–647. <http://doi.org/10.3758/s13423-015-0913-5> (URL: <http://doi.org/10.3758/s13423-015-0913-5>)
- Cumming, G. (2014). The New Statistics: Why and How. *Psychological Science*, 25(1), 7–29. <http://doi.org/10.1177/0956797613504966> (URL: <http://doi.org/10.1177/0956797613504966>)
- Drennan, J. (2003). Cognitive interviewing: Verbal data in the design and pretesting of questionnaires. *Journal of Advanced Nursing*, 42(1), 57–63. <http://doi.org/10.1046/j.1365-2648.2003.02579.x> (URL: <http://doi.org/10.1046/j.1365-2648.2003.02579.x>)
- Francis, J. J., Johnston, M., Robertson, C., Glidewell, L., Entwistle, V., Eccles, M. P., & Grimshaw, J. M. (2010). What is an adequate sample size? Operationalising data saturation for theory-based interview studies. *Psychology & Health*, 25(10), 1229–1245. <http://doi.org/10.1080/08870440903194015> (URL: <http://doi.org/10.1080/08870440903194015>)
- Haeger, H., Lambert, A. D., Kinzie, J., & Gieser, J. (2012). Cognitive Interviews to Improve Survey Instruments. Annual Forum of the Association for Institutional Research, (June), 15. Retrieved from <http://cpr.indiana.edu/uploads/AIR2012%20Cognitive%20Interviews.pdf> (URL: <http://cpr.indiana.edu/uploads/AIR2012%20Cognitive%20Interviews.pdf>)
- Hiles, D. (2012). Mixed methods and the problem of the logic of inquiry. Paper Presented at the British Psychological Society Qualitative Methods in Psychology Section Conference., 1–8. Retrieved from [http://psy.dmu.ac.uk/drhiles/pdf%27s/Hiles%20\(2012\)%20Mixed%20Methods%20\(BPS%20Conference%20-%20London\).pdf](http://psy.dmu.ac.uk/drhiles/pdf%27s/Hiles%20(2012)%20Mixed%20Methods%20(BPS%20Conference%20-%20London).pdf) (URL: [http://psy.dmu.ac.uk/drhiles/pdf%27s/Hiles%20\(2012\)%20Mixed%20Methods%20\(BPS%20Conference%20-%20London\).pdf](http://psy.dmu.ac.uk/drhiles/pdf%27s/Hiles%20(2012)%20Mixed%20Methods%20(BPS%20Conference%20-%20London).pdf))
- Leppink, J. (2017). Revisiting the quantitative–qualitative-mixed methods labels: Research questions, developments, and the need for replication. *Journal of Taibah University Medical Sciences*, 12(2), 97–101. <http://doi.org/10.1016/j.jtumed.2016.11.008> (URL: <http://doi.org/10.1016/j.jtumed.2016.11.008>)
- Mason, M. (2010). Sample Size and Saturation in PhD Studies Using Qualitative Interviews. *Forum Qualitative Sozialforschung / Forum: Qualitative Social Research*, 11(3), Art 8. Retrieved from <http://nbn-resolving.de/urn:nbn:de:0114-fqs100387> (URL: <http://nbn-resolving.de/urn:nbn:de:0114-fqs100387>)
- Maters, G. A., Sanderman, R., Kim, A. Y., & Coyne, J. C. (2013). Problems in Cross-Cultural Use of the Hospital Anxiety and Depression Scale: “No Butterflies in the Desert.” *PLoS ONE*, 8(8). <http://doi.org/10.1371/journal.pone.0070975> (URL: <http://doi.org/10.1371/journal.pone.0070975>)
- Moinester, M., & Gottfried, R. (2014). Sample size estimation for correlations with pre-specified confidence interval. *The Quantitative Methods of Psychology*, 10(2), 124–130. Retrieved from <http://tqmp.org/RegularArticles/vol10-2/p124> (URL: <http://tqmp.org/RegularArticles/vol10-2/p124>)

Peters, G.-J. Y., & Crutzen, R. (2017). Knowing exactly how effective an intervention, treatment, or manipulation is and ensuring that a study replicates: accuracy in parameter estimation as a partial solution to the replication crisis. Retrieved from <http://osf.io/preprints/psyarxiv/cjsk2> (URL: <http://osf.io/preprints/psyarxiv/cjsk2>)

Shaw, R., & Flowers, P. (2012). Overcoming our “irreconcilable epistemological differences” and moving beyond the dichotomy of qualitative versus quantitative research in health psychology. *European Health Psychologist*, 14(4), 102–103.

Viechtbauer, W., Smits, L., Kotz, D., Bude, L., Spigt, M., Serroyen, J., & Crutzen, R. (2015). A simple formula for the calculation of sample size in pilot studies. *Journal of Clinical Epidemiology*, 68(11), 1375–1379. <http://doi.org/10.1016/j.jclinepi.2015.04.014> (URL: <http://doi.org/10.1016/j.jclinepi.2015.04.014>)

Willis, G. B. (2005). *Cognitive interviewing: a tool for improving questionnaire design*. Thousand Oaks, CA, US: Sage Publishing.

Willis, G. B., & Artino, A. R. (2013). What Do Our Respondents Think We're Asking? Using Cognitive Interviewing to Improve Medical Education Surveys. *Journal of Graduate Medical Education*, 5(3), 353–6. <http://doi.org/10.4300/JGME-D-13-00154.1> (URL: <http://doi.org/10.4300/JGME-D-13-00154.1>)

Sample size tables

NHST

Correlation coefficients

Tabel 1: Power table to obtain a power of 80%, 90%, 95% en 99%, for population correlation coefficients from .1 to .9.

	80%	90%	95%	99%
.1	782	1046	1293	1828
.15	346	463	571	807
.2	194	258	319	450
.25	123	164	202	284
.3	85	112	138	194
.35	61	81	100	140
.4	46	61	75	105
.45	36	47	58	81
.5	29	38	46	64
.55	23	30	37	51
.6	19	25	30	41
.65	16	20	24	33
.7	13	17	20	27
.75	11	14	17	22
.8	9	12	14	18
.85	8	10	11	15
.9	7	8	9	12

Tabel 2: Power table for obtaining a power of 80% with in the columns population correlation coefficient values of .1 to .9, and in the rows the number of p -values computed in a study, varying from 1 to 100.

	.1	.15	.2	.25	.3	.35	.4	.45	.5	.55	.6	.65	.7	.75	.8	.85	.9
1 p-values; alpha = 0.05	782	346	194	123	85	61	46	36	29	23	19	16	13	11	9	8	7
2 p-values; alpha = 0.025	947	419	234	148	102	74	56	43	34	28	23	19	16	13	11	9	8
3 p-values; alpha = 0.0167	1043	461	258	163	112	81	61	47	38	30	25	20	17	14	12	10	8
4 p-values; alpha = 0.0125	1111	491	274	174	119	86	65	50	40	32	26	22	18	15	12	10	8
5 p-values; alpha = 0.01	1163	514	287	182	125	90	68	53	42	34	27	23	19	16	13	11	9
6 p-values; alpha = 0.0083	1206	533	298	189	129	94	71	55	43	35	28	23	19	16	13	11	9
7 p-values; alpha = 0.0071	1242	549	307	194	133	96	73	56	44	36	29	24	20	16	14	11	9
8 p-values; alpha = 0.0062	1273	563	314	199	137	99	74	58	46	37	30	25	20	17	14	12	9
9 p-values; alpha = 0.0056	1301	575	321	203	140	101	76	59	47	37	30	25	21	17	14	12	10
10 p-values; alpha = 0.005	1326	586	327	207	142	103	77	60	47	38	31	25	21	17	14	12	10
20 p-values; alpha = 0.0025	1487	657	367	232	159	115	87	67	53	42	35	28	23	19	16	13	11
30 p-values; alpha = 0.0017	1581	699	390	247	169	122	92	71	56	45	37	30	25	20	17	14	11
40 p-values; alpha = 0.0012	1648	728	406	257	176	127	96	74	58	47	38	31	26	21	17	14	11
50 p-values; alpha = 0.001	1700	751	419	265	182	131	99	76	60	48	39	32	26	22	18	15	12
100 p-values; alpha = 0.0005	1860	822	458	290	199	144	108	83	66	53	43	35	29	24	19	16	13

Tabel 3: Power table for obtaining a power of 95% with in the columns population correlation coefficient values of .1 to .9, and in the rows the number of *p*-values computed in a study, varying from 1 to 100.

	.1	.15	.2	.25	.3	.35	.4	.45	.5	.55	.6	.65	.7	.75	.8	.85	.9
1 <i>p</i>-values; alpha = 0.05	1293	571	319	202	138	100	75	58	46	37	30	24	20	17	14	11	9
2 <i>p</i>-values; alpha = 0.025	1503	664	370	234	160	116	87	67	53	42	34	28	23	19	16	13	10
3 <i>p</i>-values; alpha = 0.0167	1623	717	400	253	173	125	94	72	57	46	37	30	25	20	17	14	11
4 <i>p</i>-values; alpha = 0.0125	1707	754	420	266	182	131	99	76	60	48	39	32	26	21	17	14	11
5 <i>p</i>-values; alpha = 0.01	1772	783	436	276	189	136	102	79	62	50	40	33	27	22	18	15	12
6 <i>p</i>-values; alpha = 0.0083	1825	806	449	284	194	140	105	81	64	51	41	34	28	23	18	15	12
7 <i>p</i>-values; alpha = 0.0071	1870	826	460	291	199	144	108	83	65	52	42	34	28	23	19	15	12
8 <i>p</i>-values; alpha = 0.0062	1908	842	470	297	203	147	110	85	67	53	43	35	29	23	19	15	12
9 <i>p</i>-values; alpha = 0.0056	1942	857	478	302	207	149	112	86	68	54	44	36	29	24	19	16	12
10 <i>p</i>-values; alpha = 0.005	1972	871	485	307	210	151	114	87	69	55	44	36	30	24	20	16	13
20 <i>p</i>-values; alpha = 0.0025	2168	957	533	337	231	166	125	96	75	60	49	40	32	26	22	17	14
30 <i>p</i>-values; alpha = 0.0017	2281	1007	561	355	243	175	131	101	79	63	51	42	34	28	23	18	14
40 <i>p</i>-values; alpha = 0.0012	2361	1042	581	367	251	181	136	104	82	66	53	43	35	29	23	19	15
50 <i>p</i>-values; alpha = 0.001	2423	1070	596	377	258	186	139	107	84	67	54	44	36	29	24	19	15
100 <i>p</i>-values; alpha = 0.0005	2613	1154	643	406	278	200	150	115	91	72	58	47	39	32	26	21	16

Cohen's *d*

Tabel 4: Power table to obtain a power of 80%, 90%, 95% en 99%, for population Cohen's d values from .1 to 1.3.

	80%	90%	95%	99%
0.1	3142	4206	5200	7352
0.2	788	1054	1302	1840
0.3	352	470	580	820
0.4	200	266	328	462
0.5	128	172	210	296
0.6	90	120	148	208
0.7	68	88	110	152
0.8	52	68	84	118
0.9	42	54	68	94
1	34	46	54	76
1.1	30	38	46	64
1.2	24	32	40	54
1.3	22	28	34	46

Tabel 5: Power table for obtaining a power of 80% with in the columns population Cohen's d values of .1 to 1.3, and in the rows the number of p -values computed in a study, varying from 1 to 100.

	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	1.1	1.2	1.3
1 p-values; alpha = 0.05	3142	788	352	200	128	90	68	52	42	34	30	24	22
2 p-values; alpha = 0.025	3806	954	426	242	156	110	82	62	50	42	36	30	26
3 p-values; alpha = 0.0167	4192	1050	470	266	172	120	90	70	56	46	38	34	28
4 p-values; alpha = 0.0125	4464	1120	500	282	182	128	96	74	60	48	42	36	30
5 p-values; alpha = 0.01	4676	1172	524	296	192	134	100	78	62	52	44	36	32
6 p-values; alpha = 0.0083	4848	1216	542	308	198	140	104	80	64	54	44	38	34
7 p-values; alpha = 0.0071	4994	1252	558	316	204	144	106	82	66	54	46	40	34
8 p-values; alpha = 0.0062	5120	1284	574	324	210	146	110	84	68	56	48	40	36
9 p-values; alpha = 0.0056	5230	1312	586	332	214	150	112	86	70	58	48	42	36
10 p-values; alpha = 0.005	5330	1336	596	338	218	152	114	88	70	58	50	42	36
20 p-values; alpha = 0.0025	5980	1500	670	380	244	172	128	98	80	66	56	48	42
30 p-values; alpha = 0.0017	6360	1594	712	404	260	182	136	106	84	70	58	50	44
40 p-values; alpha = 0.0012	6628	1662	742	420	272	190	142	110	88	72	62	52	46
50 p-values; alpha = 0.001	6836	1714	766	434	280	196	146	114	90	74	62	54	46
100 p-values; alpha = 0.0005	7480	1876	838	474	306	214	160	124	100	82	68	60	52

Tabel 6: Power table for obtaining a power of 95% with in the columns population Cohen's d values of .1 to 1.3, and in the rows the number of p -values computed in a study, varying from 1 to 100.

	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	1.1	1.2	1.3
1 p-values; alpha = 0.05	5200	1302	580	328	210	148	110	84	68	54	46	40	34
2 p-values; alpha = 0.025	6044	1514	674	382	246	172	126	98	78	64	54	46	40
3 p-values; alpha = 0.0167	6528	1636	728	412	264	186	138	106	84	70	58	50	42
4 p-values; alpha = 0.0125	6868	1720	766	434	278	194	144	112	88	72	60	52	44
5 p-values; alpha = 0.01	7130	1786	796	450	290	202	150	116	92	76	64	54	46
6 p-values; alpha = 0.0083	7342	1838	820	464	298	208	154	120	96	78	66	56	48
7 p-values; alpha = 0.0071	7522	1884	840	474	306	214	158	122	98	80	66	56	50
8 p-values; alpha = 0.0062	7676	1922	858	484	312	218	162	124	100	82	68	58	50
9 p-values; alpha = 0.0056	7812	1956	872	492	318	222	164	126	102	82	70	60	52
10 p-values; alpha = 0.005	7932	1986	886	500	322	226	166	128	102	84	70	60	52
20 p-values; alpha = 0.0025	8722	2184	974	550	354	248	184	142	114	92	78	66	58
30 p-values; alpha = 0.0017	9180	2300	1026	580	372	260	194	150	120	98	82	70	60
40 p-values; alpha = 0.0012	9502	2380	1062	600	386	270	200	154	124	102	84	72	62
50 p-values; alpha = 0.001	9750	2442	1088	616	396	278	206	158	126	104	88	74	64
100 p-values; alpha = 0.0005	10516	2634	1174	664	428	298	222	172	136	112	94	80	70

Parameter estimation

Correlation coefficients

Tabel 7: Sample size table for obtaining a correlation coefficient estimate with a confidence interval with 95% confidence for population correlation coefficient values from .1 to .9 and desired half-widths (margins of error) from 0.05 to 0.5.

	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
$r = 0.1$	1507	378	168	95	61	43	32	25	20	16
$r = 0.2$	1417	355	159	90	58	41	30	24	19	16
$r = 0.3$	1274	320	143	81	53	37	28	22	18	15
$r = 0.4$	1086	273	123	70	46	32	24	19	16	13
$r = 0.5$	867	219	99	57	37	27	20	16	13	11
$r = 0.6$	633	161	74	43	29	21	16	13	11	10
$r = 0.7$	404	105	49	30	21	15	12	10	9	8
$r = 0.8$	205	56	28	18	14	11	9	8	7	7
$r = 0.9$	63	21	13	11	9	8	7	7	6	6

Tabel 8: Sample size table for obtaining a correlation coefficient estimate with a confidence interval with 99% confidence for population correlation coefficient values from .1 to .9 and desired half-widths (margins of error) from 0.05 to 0.5.

	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
$r = 0.1$	2600	650	288	162	103	72	53	40	32	26
$r = 0.2$	2446	611	271	153	98	68	50	38	30	25
$r = 0.3$	2198	550	245	138	88	61	45	35	28	23
$r = 0.4$	1874	469	209	118	76	53	40	31	24	20
$r = 0.5$	1495	376	168	96	62	44	33	26	21	17
$r = 0.6$	1091	276	125	72	47	34	26	20	17	14
$r = 0.7$	696	178	82	48	33	24	19	16	13	11
$r = 0.8$	352	94	46	29	21	16	13	12	10	9
$r = 0.9$	106	34	21	17	14	12	11	9	8	7

Cohen's d

Tabel 9: Sample size table for obtaining a Cohen's d estimate with a confidence interval with 95% confidence for population values of Cohen's d of 0.2, 0.5 and 0.8, and desired half-widths (margins of error) from 0.05 to 0.5.

	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
$d = 0.2$	6161	1537	681	383	245	170	125	95	75	61
$d = 0.5$	6296	1564	691	386	246	170	124	95	74	60
$d = 0.8$	6570	1626	716	399	253	174	127	96	75	60

Tabel 10: Sample size table for obtaining a Cohen's d estimate with a confidence interval with 99% confidence for population values of Cohen's d of 0.2, 0.5 and 0.8, and desired half-widths (margins of error) from 0.05 to 0.5.

	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5
$d = 0.2$	10643	2655	1178	662	423	294	216	165	131	106
$d = 0.5$	10879	2704	1195	669	426	295	216	164	130	105
$d = 0.8$	11354	2812	1239	691	438	302	220	167	131	106