Using simulation and resampling to improve the
statistical power and reproducibility of psychological research

Roger W. Strong & George A. Alvarez

Department of Psychology

Harvard University

Cambridge, MA, USA

Abstract

The replication crisis has brought about an increased focus on improving the reproducibility of psychological research (Open Science Collaboration, 2015). Although some failed replications reflect false-positives in original research findings, many are likely the result of low statistical power, which can cause failed replications even when an effect is real, no questionable research practices are used, and an experiment's methodology is repeated perfectly. The present paper describes a simulation method (bootstrap resampling) that can be used in combination with pilot data or synthetic data to produce highly powered experimental designs. Unlike other commonly used power analysis approaches (e.g., G*Power), bootstrap resampling can be adapted to any experimental design to account for various factors that influence statistical power, including sample size, number of trials per condition, and participant exclusion criteria. Ignoring some of these factors (e.g., by using G*Power) can overestimate the power of a study or replication, increasing the likelihood that your findings will not replicate. By demonstrating how these factors influence the consistency of experimental findings, this paper provides examples of how simulation can be used to improve statistical power and reproducibility. Further, we provide a MATLAB toolbox that can be used to implement these simulation-based methods on existing pilot data (<https://harvard-visionlab.github.io/power-sim>).

Statistical Power in Psychology

Reproducing previous experimental results is critical for scientific advancement, yet many psychological findings are difficult to replicate (Open Science Collaboration, 2015). Although some of these failed replications reflect false-positives caused by “questionable research practices” (Gelman & Loken, 2013; John, Loewenstein, & Prelec, 2012; Simmons, Nelson, & Simonsohn, 2011), many failed replications likely result from low statistical power, the likelihood of observing a statistically significant result when an effect is present (i.e., the probability of rejecting the null hypothesis when the null hypothesis is false). In addition to being less likely to produce statistically significant results, studies with low power are more likely to produce largely exaggerated effect sizes by chance (Button et al., 2013). As demonstrated with examples in this paper, overestimation of effect size due to low statistical power can lead to replication failures, even when the direction of a reported effect is real, no questionable research practices are used, and an experiment’s methodology is perfectly repeated. Moreover, we highlight that using standard power estimates (i.e., G*Power; Erdfelder, Faul, & Buchner, 1996) that ignore measurement variability can overestimate power — sometimes dramatically — increasing the likelihood that other researchers fail to replicate your work when relying on the same overestimate of power, even when the effects you report are real.

Statistical power influences reproducibility because of the variability inherent to experimental outcomes – no two experiments produce identical results, even when sampling from the same population with identical methodology. This variance in experimental outcomes occurs for two reasons: sampling variability and measurement variability. *Sampling variability* occurs because experiments sample a random subset of a population (as opposed to the entire population), with this subset varying across individual experiments. For example, imagine twice

sampling the height of 20 random individuals on a college campus. Despite sampling from the same population (students at the college) both times, the two rounds of data collection would sample different individuals, and thus likely would produce a different average height. Similarly, a psychological study may fail to replicate because of sampling variability between the original study and the replication.

Whereas sampling variability occurs because experiments cannot sample the entire population, *measurement variability* occurs because of measurement imprecision inherent to most experimental designs. Importantly, this imprecision is not usually the result of experimenter or technical error (e.g., a faulty keyboard), but instead due to the probabilistic nature of most experimental measurements. For example, imagine investigating whether people have the ability to predict the future by asking them to twice predict the outcome of 10 flips of a fair coin (which has a 50% chance of landing heads and a 50% chance of landing tails). Even if none of the participants had the ability to predict the future (and thus all had a 50% chance of correctly guessing each individual flip), only about 6% of participants would guess exactly 5 out of 10 flips correctly both rounds. Similarly, studies measuring reaction time, accuracy, ratings, or nearly any other measure will have some degree of measurement variability – in reaction time studies, for example, participants will not have the same exact reaction time for each trial, even if the trials are identical. Because of measurement variability, even the same individual is unlikely to produce identical results when completing the same test multiple times (even if there were no practice effects).

Simulating Variability with Bootstrap Resampling

Creating experiments with high statistical power requires researchers to minimize both the sampling variability and measurement variability of their experimental designs, although measurement variability tends to be ignored in standard power estimates (e.g., G*Power). This paper uses the Monte Carlo simulation technique of bootstrap resampling (also known as bootstrapping; Efron, 1979) to demonstrate how various features of experimental design (e.g., sample size, number of trials, and participant exclusion criteria) influence sampling and measurement variability, and thus experimental power. The bootstrap technique allows researchers to simulate running an experimental design many times, with each round of simulated data collection sampling from a data set representative of the population of interest. For example, a pilot set of data can be used to approximate the population a main experiment would sample from – this approach will be used in this paper’s examples. To our knowledge, using simulation techniques to determine experimental design is rare in psychological science, despite being used more often in other scientific fields (e.g., Green & MacLeod, 2015; Johnson, Barry, Ferguson, & Müller, 2015; Muthén & Muthén, 2002) – thus, in the present paper we provide an overview of how bootstrap resampling can be used to create highly powered and reproducible experimental designs.

Bootstrapping allows estimation of sampling variability by sampling participants randomly with replacement during each simulated experiment. Note that individual participants are sampled a variable number of times each simulated experiment (e.g., during some simulated experiments a particular participant may not be sampled at all, while during other simulated experiments the same participant will be sampled multiple times; see Figure 1). This procedure allows a design’s sampling variability to be estimated from a single set of data. Furthermore,

Figure 1: Simulation via bootstrap resampling. With bootstrapping, a data set (e.g., pilot data) is used to represent a population of interest. Experimental designs can be simulated from this data set by repeatedly sampling participants (represented by different colors) and those participants’ individual trials (represented by different shades of each color) with replacement. Each simulated experiment is unique for two reasons: 1) each simulated experiment will have different participants (due to resampling participants from the pilot data with replacement – see variation in colors in simulated data), and 2) each participant can have a different score each time they are sampled (due to resampling their trials with replacement – note differences between the first and fourth row of the first simulated data set, where different trials of the same participant have been randomly sampled). Although in this example the simulated data sets have the same structure as the pilot data (9 participants and 9 trials), bootstrap resampling allows the simulation of experimental designs that differ from the pilot study.

sampling participants with replacement enables the simulation of experimental designs across different sample sizes (for example, a pilot data set of 30 participants can be used to simulate an experiment with 100 – or any number – of participants).

Whereas sampling variability can be estimated by resampling participants with replacement, a design’s measurement variability can be estimated by resampling each participant’s individual trials with replacement (e.g., during some simulated experiments a particular trial may not be sampled at all, while during other simulated experiments the same trial will be sampled multiple times; see Figure 1). Using a single data set, this procedure allows estimation of how the performance of each individual participant will vary across repetitions of

the same experimental design, while additionally allowing simulation of experimental designs with varying amounts of trials. In summary, resampling participants with replacement allows estimation of a design's sampling variability, while resampling each participant's individual trials with replacement allows estimation of a design's measurement variability.

This paper provides examples of bootstrap resampling using a synthetic data set generated by the authors for demonstration. For the sake of clarity, we will refer to this synthetic data set as “pilot data” (and its individual participants as “pilot participants”) collected from a “pilot study,” as we recommend collecting pilot data for power simulations in actual experiments (rather than creating synthetic data that you hope approximates the population you will sample from; see discussion section). This pilot experiment used a within-subjects design, requiring observers ($N = 30$) to complete 32 trials for each of two experimental conditions, on a task that required observers to make a two-alternative forced-choice response (resulting in chance performance being 50%). The pilot participants showed better performance for the second condition ($M = 84.3\%$, $SD = 15.8\%$) than for the first condition ($M = 77.4\%$, $SD = 12.7\%$), a statistically significant result ($t(29) = 3.55$, $p < .001$, $d_z = .6475$).

The remainder of this paper will demonstrate how researchers could use this pilot set of data and bootstrap resampling to create a highly powered main experiment. For each experimental design (defined by a combination of sample size N and trials per condition T) for which we wished to calculate statistical power, 100,000 experiments were simulated by repeating the following procedure 100,000 times (simulating the 40 designs presented in Figure 4 takes about 30 seconds total in MATLAB):

1. Randomly select N participants from the pilot data set, with replacement.
2. Each time a participant is selected, randomly select T of that participant's trials from each of the two conditions, with replacement.

3. For each of the N participants, calculate the mean score for each of the two conditions across the T trials per condition.
4. For each of the N participants, calculate the difference between the mean score of each of the two conditions.
5. Calculate the study's effect size d_z (for a review of different measures of effect size, see Lakens, 2013) by dividing the mean of the N difference scores by the standard deviation of the N difference scores.
6. Determine whether the study reaches statistical significance (using a two-tailed, paired-samples t-test with $\alpha = .05$) in the predicted direction.

After completing this procedure 100,000 times, the power of the simulated experimental design can be calculated using the percentage of simulated studies that reach statistical significance in the predicted direction (e.g., if 90,000 simulated experiments reached significance in the predicted direction, the power of the simulated experimental design would be .90).

Additionally, we can examine the distribution of effect sizes obtained from the simulated design, allowing calculation of the design's expected effect size (i.e., the average observed effect size) and the variability of observed effect sizes. Importantly, although our pilot data set had 30 participants with 32 trials per condition, bootstrap resampling allows simulation of designs with any number of participants and trials per condition (due to sampling randomly with replacement). Additionally, although this paper uses the example of a within-subjects study with two conditions, the bootstrap resampling procedure can be modified for any study design or analysis plan by following the steps of: 1) collecting pilot data, 2) resampling from that pilot data to simulate running candidate experimental designs, and 3) calculating how often those designs produce statistically significant results using your planned analyses.

Using Bootstrap Resampling to Maximize Power

How Sample Size Influences Power

Perhaps the most widely used strategy for maximizing statistical power is using a large sample size. For any given sample size, there is a minimum effect size necessary for obtaining a statistically significant result. For example, the critical effect size for a study of sample size $N = 30$ is $d_z = 0.3734$ (when comparing two within-subjects conditions using a two-tailed t-test with $\alpha = .05$) – studies with effect sizes less than this critical value will have p-values above .05, and studies with effect sizes greater than this critical value will have p-values below .05. As sample size increases, the critical effect size necessary for declaring results statistically significant decreases; in other words, increasing sample size increases statistical power by allowing smaller effect sizes to be declared statistically significant. Importantly, increasing sample size does *not* increase the expected effect size of an experimental design (in fact, the expected effect size drops slightly when increasing sample size, as small samples systematically underestimate the variance of their underlying population; Holtzman, 1950). This fact is evident in Figure 2, which shows the distribution of effect sizes at various sample sizes (obtained by bootstrap resampling our pilot data set at different values of sample size N , with trials per condition $T = 32$). The average effect size (solid black line) does not change much with increasing sample size. Instead, the critical effect size necessary for declaring results significant (black circle at intersection of green and red) decreases, resulting in higher power (area under the green portion of the distribution) with larger sample sizes.

In addition to lowering the critical effect size necessary for declaring results significant, larger sample sizes reduce the variance of observed effect sizes, due to reduced sampling variability (note the smaller standard deviations of effect size distributions as sample size increases in Figure 2). Smaller effect size variance increases statistical power because of reduced probability of observing a large underestimation of effect size (as the tails of the distribution

Figure 2: Effect size (d_z) distributions at increasing sample sizes (N), while holding number of trials constant (32 trials per condition, 100,000 simulations per panel). Although the average observed effect size (solid black line) does not vary much with increasing sample size, power (area under green portion of curve) increases with increasing sample size. This is because increasing sample size both 1) decreases the critical effect size need to achieve statistical significance (black circle at intersection of red and green portion of curve), and 2) decreases the variance in observed effect size (dashed black lines represent ± 2 SD from the mean observed effect size). Note that these distributions and power calculations are specific to the pilot data used for resampling.

move closer to the average effect size). In addition to increasing statistical power, reduced effect size variance increases the probability of measuring an effect size close to a design's expected effect size, thus increasing the likelihood of a study replicating successfully. With small sample sizes, high sampling variability causes the variance of observed effect sizes to be very large; this often produces a situation where only effect sizes larger than the design's expected effect size produce statistical significance (e.g., see Figure 2 bottom panel where $N = 12$); therefore, small sample studies can only produce statistically significant results when a very large effect size is observed by chance (unless their expected effect size is very large). When other researchers attempt to replicate a reported effect size that is larger than a design's expected effect size, they will base their sample size on the overestimated effect size, making replication of the original finding very unlikely (Button et al., 2003). In summary, minimizing sampling variability by running experiments with a large sample size is important for two reasons: 1) increasing the probability of obtaining statistically significant results, and 2) increasing the probability that observed data is a good estimate of an experiment's expected effect size, thus increasing the probability of the study being replicated successfully.

How Number of Trials Influences Power

Whereas large sample sizes increase statistical power by both decreasing sampling variability and reducing the critical effect size necessary for declaring statistical significance, increasing the number of trials per condition can increase a study design's power by reducing measurement variability. Reducing measurement variability increases a study's expected effect size, the ratio of an observed effect's magnitude (e.g., difference between conditions) over its variability. All else being equal, increasing the number of trials per condition produces a more reliable estimate

of an experiment's dependent measures, reducing the variance of the observed effect (the denominator of the ratio determining effect size). For example, consider 100,000 people each flipping a fair coin 10 times. Although when averaging across individuals the mean percentage of "heads" flipped will be 50%, there will be variability in this percentage ($SD = 15.8\%$). If the same 100,000 people flipped a coin 100 times each, the average percentage of heads would again be 50%, but this time with much less variability ($SD = 5.0\%$). As the number of flips approaches infinity, the variability in the percentage of heads would approach zero, as each individual would consistently produce data reflecting the true underlying value of what is being measured (a fair coin that has a 50% chance of landing heads). Similarly, increasing the number of trials in an experimental design increases the likelihood that data are representative of each participant's true underlying ability (assuming that their ability does not change as the number of trials changes), increasing the power and reproducibility of the study design.

Figure 3 demonstrates (via bootstrapping our pilot data) how decreasing measurement variability (by increasing the number of trials per condition) increases expected effect size, and thus power. While keeping sample size N constant, note that increasing number of trials per condition T shifts the distribution of observed effect sizes rightward (while the critical effect size remains constant), increasing experimental power. Importantly, increasing the number of trials per condition will only decrease measurement variability (and thus increase power) if observers maintain the same underlying performance level that they had with fewer trials. Factors such as fatigue, boredom, or practice effects could contribute to measurement variability with large trial amounts, potentially outweighing the benefits of more trials. Therefore, experimenters should select a trial amount that is large enough to minimize measurement variability, without being likely to cause participants to become differentially engaged with the task.

Figure 3: Effect size (d_z) distributions at increasing numbers of trials per condition (T). When holding sample size constant ($N = 36$, 100,000 simulations per panel), more trials results in a greater average observed effect size (solid black line), resulting in greater power (area under green portion of curve). Note that increasing number of trials does not cause a large change in the variance of observed effect sizes (dashed black lines represent ± 2 SD from the mean observed effect size). Additionally, the critical effect size (black circle at intersection of red and green portion of curve) remains the same when sample size is held constant. Note that these distributions and power calculations are specific to the pilot data used for resampling.

Selecting an Optimal Sample Size and Number of Trials

The previous sections used bootstrap resampling to demonstrate how sample size (sampling variance) and trials per condition (measurement variance) separately influence statistical power. Although the previous sections held either sample size or trials per condition constant while varying the other factor, experimenters can use bootstrap resampling to explore

Figure 4: Power to detect the effect present in our pilot data as a function of sample size and number of trials per condition. By resampling their pilot data, experimenters can determine the best design for optimizing power. After bootstrapping the pilot data used in this example, if an experimenter desired a power level of .95, they could choose from running 60 participants with 32 trials per condition, 72 participants with 24 trials per condition, or 84 participants with 16 trials per condition. Note that these statistical power estimates are specific to the pilot data used for resampling.

how different combinations of sample size and trials per condition influence the power of their experiment. Figure 4 depicts statistical power estimates for our hypothetical study (generated via bootstrap resampling our pilot data) for experimental designs across various combinations of trials per condition and sample size. Using this output, a researcher wishing to run a main experiment with a statistical power of .95 could choose a design with 60 participants and 32 trials per condition, 72 participants with 24 trials per condition, or 84 participants with 16 trials per condition (note that these power estimates are specific to the pilot data set used in this paper). Tools and tutorials for bootstrapping your own pilot data to maximize the power of your experimental designs is available at <https://harvard-visionlab.github.io/power-sim>. Although the example in this paper is a simple within-subjects design comparing two conditions, bootstrap resampling can be used to estimate the power of any experimental design.

Determining Exclusion Criteria

Participants performing at ceiling (due to the task being too easy) or floor (due to the task being too hard, or participants not completing the task correctly) decrease the expected effect size of a study by both A) decreasing the magnitude of the observed effect and B) increasing the variability of the observed effect. Participants at ceiling will score at or near 100% in each condition, resulting in a difference of 0 being averaged into the estimate of the magnitude of difference between conditions. For these participants, any differences between conditions will likely be due to lapses of attention, and not an underlying effect of interest, increasing the variability of the observed effect as well. The same problem occurs for participants at floor, who perform at or near chance levels. On average, participants at floor will show no difference

between conditions, while also having great variability in the observed difference between conditions, especially when study designs have a small number of trials per condition.

Ceiling and floor effects can be avoided in a number of ways. One strategy is using a thresholding procedure to choose a difficulty level that is specific to each observer (e.g., Watson & Pelli, 1983). While this can be helpful for observers who are engaged with the task, this approach will not prevent observers who are simply guessing or performing the task incorrectly from performing at chance levels. Although these individuals can sometimes be identified using attention checks (Oppenheimer, Meyvis, & Davidenko, 2009), another way to exclude participants who are completely guessing is using a performance cutoff. Figure 5A and Figure 5B demonstrate how researchers might go about selecting exclusion criteria (either by simulating the performance of randomly guessing participants, or by using a binomial calculator). Consider our pilot study with 64 total trials (32 trials per condition), where chance-level performance was 50% (due to its two-alternative forced-choice design). Although choosing a 50% performance cutoff (meaning that participants' overall score must be above 50% to be included in the final analysis) may seem intuitive, using this cutoff would result in only 55% of randomly guessing participants being excluded (Figure 5A, magenta circle). By instead using a cutoff of 60% (blue circle), 95% of randomly guessing participants would be excluded. For any selected exclusion level, researchers should also consider the exclusion rate of participants whose true ability is above chance level (i.e., participants who are not completely guessing). Figure 5B shows the expected exclusion rate of our design as a function of true participant ability (at three different performance cutoffs). As the accuracy cutoff increases, the exclusion rate of both randomly guessing participants (50% true ability) and non-guessing participants ($> 50\%$ true ability) increases; thus, the benefit of excluding more randomly guessing participants comes at the cost

Figure 5: Selection and impact of participant exclusion criteria. **A:** Percentage of randomly guessing participants expected to be excluded at various accuracy cutoffs (for a 2 AFC experiment with 64 total trials; magenta circle = 50% accuracy cutoff, blue circle = 60% accuracy cutoff, cyan circle = 70% accuracy cutoff). **B:** Expected exclusion percentage versus true ability (separate for each potential accuracy cutoff). The higher the accuracy cutoff, the larger the percentage of guessing participants (50% true ability) that will be excluded, but more non-guessing participants (> 50% true ability) that will be excluded as well. **C:** Distribution of effect sizes (d_z) with (top panel) and without (bottom panel) using a 60% accuracy cutoff for exclusion, generated using bootstrap resampling ($N = 36$, $T = 32$, excluded participants replaced).

of excluding more non-guessing participants. Fortunately, the effect of excluding participants on any specific design can be estimated by bootstrap resampling pilot data.

To demonstrate the potential utility of using exclusion criteria, we bootstrapped our pilot data with an additional simulation parameter – requiring simulated participants to have greater than 60% performance (aggregated across conditions) to be included in the final analysis. Importantly, this cutoff was applied to simulated participants, and was NOT applied to pilot participants – all pilot participants were eligible to be sampled, no matter their score. For example, pilot participants with 50% accuracy could score above the 60% cutoff when their trials were resampled as simulated participants. This approach maintains the pilot data set as a representative example of the population, allowing estimation of how many participants will be excluded when using the selected exclusion criteria, as well as assessment of whether including exclusion criteria will increase statistical power. Note that for our simulation using exclusion criteria, any excluded participants were replaced until the target sample size was reached. As depicted in Figure 5C, excluding and replacing participants with accuracy 60% or below (6.3 participants excluded and replaced, on average) results in 98% power, compared to 81% power when no exclusion occurred (with $N = 36$ and $T = 32$ in both cases); thus, including exclusion criteria would increase the statistical power of our main experiment.

Although the previous example used accuracy as the parameter for determining exclusion, other forms of exclusion criteria can be used as well. For example, a reaction time study could simulate the impact of excluding participants whose overall reaction time was outside a specified range (e.g., 300 ms – 2000 ms), or the impact of excluding individual trials outside of the range. For studies measuring participants' ratings, reliability-based exclusion cutoffs can be simulated. For example, in a study examining participants' ratings of the

attractiveness of faces, experimenters could present each face twice, allowing calculation of the reliability of each participant's ratings. Participants assigning ratings arbitrarily would be likely to have low reliability in their ratings; experimenters could use bootstrap resampling on a set of pilot data to determine a reliability cutoff that maximizes the power of their main experiment.

Importantly, choices about exclusion criteria must be decided upon (and ideally should be preregistered) before running a main study for publication. Using simulation to explore different analysis plans on pilot data is an effective way to increase statistical power without inflating false-alarm rates, but choosing exclusion criteria after data for publication has been collected will lead to increased false-positives (Simmons et al., 2011). Additionally, exclusion criteria must be applied either to data aggregated across conditions (as done in the example depicted in Figure 5C), or uniformly to each separate condition, so not to bias the experimental results in any particular direction (e.g., it would be inappropriate to apply the exclusion criteria only to the condition you predict participants to do best on).

Comparing Bootstrap Resampling to Other Approaches

*G*Power*

One of the most widely used tools for choosing sample size is G*Power (Erdfelder et al., 1996; Faul, Erdfelder, Lang, & Buchner, 2009), which allows experimenters to select their sample size by entering an expected effect size and desired power level. Although G*Power accounts for sampling variability in its power estimates, it does not account for measurement variability, instead assuming no measurement variability is present. Although studies with low measurement variability will have power levels similar to G*Power's estimates, studies with even moderate measurement variability may have much lower power than estimated by

G*Power. To demonstrate, imagine determining the design of a main experiment using the pilot data discussed in this paper ($N = 30$, $T = 32$). When using G*Power to run an a priori power analysis for testing the difference between two dependent means (matched pairs) with a t-test, G*Power indicates that 33 participants are needed to obtain 95% power with an effect size $d_z = .6475$ (the effect size observed in the pilot data) and $\alpha = .05$, without considering the number of trials per condition. However, simulation using bootstrap resampling (which accounts for measurement variability by resampling trials) indicates that this experimental design suggested by G*Power ($N = 33$, $T = 32$) has only a 77% chance of detecting the effect present in the pilot data. To achieve 95% power with 32 trials per condition, bootstrap resampling estimates that 60 participants would need to be sampled (see Figure 6). This difference between expected (95%) versus actual (77%) power grows more problematic with increasing numbers of experiments; for example, a researcher needing three experiments to reach statistical significance to publish a paper would have a $.77^3 = 46\%$ chance of all three experiments reaching statistical significance when using the design suggested by G*Power, versus a $.95^3 = 86\%$ chance of all three experiments reaching statistical significance when using the design suggested by bootstrap resampling. Thus, although G*Power is an extremely useful tool (and far better than conducting no power analysis), its estimates of power are often exaggerated due to failing to account for measurement variability.

Figure 6: Comparison of effect size distributions for designs suggested by bootstrapping (top panel) and G*Power (bottom panel) for achieving 95% power to detect the effect present in the example pilot data ($N = 30$, $T = 32$ trials per condition). Simulation using bootstrap sampling, which accounts for both sampling and measurement variability, estimates that the design suggested by simulation ($N = 60$, $T = 32$) has the desired 95% power. Simulating the design suggested by G*Power ($N = 33$, $T = 32$) estimates that this design would only have 77% power, due to G*Power only accounting for sampling variability (and not measurement variability) in its power calculations.

Using the Design of Similar Studies

Another common strategy for determining experimental design is using the parameters (e.g., sample size, number of trials per condition, exclusion criteria) of previous studies exploring similar effects. Although this approach may be better than arbitrarily selecting the design of an experiment, even small changes to the sampled population or to experimental design can alter the strength of an effect. Additionally, just because a previous study found a statistically significant effect does not necessarily mean that study design had good statistical power (Anderson, Kelley, & Maxwell, 2017). Therefore, collecting pilot data is advisable when possible, both to explore

whether the magnitude an effect is similar to what has previously been reported, and to allow for power analyses via simulation.

Simulating From Population Models

When collecting pilot data is not possible, experimenters can attempt to model the population they are sampling, allowing simulation of various experimental designs from a theoretical population rather than from pilot data (Muthén & Muthén, 2002). Although accurately determining the parameters of a theoretical population may be challenging, this approach offers the same flexibility as bootstrap resampling from pilot data. Therefore, this approach has the advantage of not needing pilot data for determining experimental design, but the disadvantage of somewhat speculatively determining the characteristics of a population of interest, rather than estimating those characteristics from pilot data.

General Discussion

How Much Power Is Sufficient?

A general rule of thumb for acceptable statistical power is .80 (Cohen, 1988). While this may be sufficient for an individual study, researchers should strive for greater power when possible, particularly when running multiple experiments. For example, again consider a researcher hoping to publish a paper with three experiments. If each of those experiments had a power of .80, the likelihood of obtaining statistically significant results for all three experiments is $.80^3 = .51$. Investing great time and resources into running three experiments for the reward of a coin-flip does not seem like an advisable strategy. By designing each individual experiment to have a power of .95, the same researcher would have a much better chance of observing a

statistically significant result in all three experiments ($.95^3 = .86$). As evident from a paper finding that many studies have less than a 10% chance of all their experiments reaching statistical significance (Francis, 2014), experimenters need designs with very high power when conducting multiple experiments; one way to achieve this is by using simulation to account for both sampling and measurement variability when determining experimental design.

Strengths and Limitations of Simulation

The major strength of using simulation to determine experimental design is its large adaptability; any study design can be modeled, allowing researchers to estimate how sample size, number of trials, exclusion criteria, or any other design or analysis choice will impact their statistical power. Compared to techniques that only consider sampling variability, simulation allows researchers to produce experimental designs that are efficient, highly powered, and more likely to be reproducible.

One limitation of using bootstrap resampling is the need to collect pilot data (or generate synthetic data from estimates of population parameters). When running studies that are expensive and time consuming, collecting a round of pilot data to use for simulation may appear daunting. In some situations, therefore, researchers may be limited to basing their design off previous studies, or simulating based upon synthetic data generated from theoretical population parameters (Muthén & Muthén, 2002). However, for researchers with the resources to collect pilot data sets (particularly those who use platforms like Amazon Mechanical Turk), the investment of collecting a pilot data set should be worth the reward of creating a highly powered and reproducible experiment.

Another limitation of bootstrapping resampling is its dependence on quality pilot data. If the pilot data that has been collected is not representative of the population that will be sampled in the main experiment, then power estimates will be inaccurate. However, this is not a limitation that is specific to bootstrap resampling – an inaccurate effect size generated by unreliable pilot data will also be misleading when plugged into G*Power. Therefore, when using bootstrap resampling or any other power calculation technique, researchers can maximize the utility of their power analysis by collecting a reasonably large pilot data set, maximizing the likelihood that the pilot data is representative of the population being sampled (for guidelines on determining the sample size of a pilot study, see Hertzog, 2008).

Conclusions

Even in carefully controlled experimental designs, sampling variability and measurement variability will cause variation in the strength of an observed effect. By bootstrap resampling pilot data, experimenters can select design parameters (e.g., sample size, number of trials, and exclusion criteria) that minimize these two forms of variability, and thus increase the statistical power and reproducibility of their work. To facilitate researchers' use of these simulation-based power analyses, we have made a toolbox available at <https://harvard-visionlab.github.io/power-sim>.

Author Contributions

R. W. Strong and G. A. Alvarez jointly contributed to writing this manuscript. R. W. Strong created the Power-Sim Toolbox.

Declaration of Conflicting Interests

The authors declare no conflicts of interest with respect to the authorship or the publication of this article.

References

- Anderson, S. F., Kelley, K., & Maxwell, S. E. (2017). Sample-size planning for more accurate statistical power: A method adjusting sample effect sizes for publication bias and uncertainty. *Psychological Science*, 28(11), 1547-1562.
- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365-376.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.
- Efron, B. (1979). Bootstrap methods: another look at the jackknife. *The Annals of Statistics*, 7(1), 1-26.
- Erdfelder, E., Faul, F., & Buchner, A. (1996). GPOWER: A general power analysis program. *Behavior Research Methods, Instruments, & Computers*, 28(1), 1-11.
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41(4), 1149-1160.
- Francis, G. (2014). The frequency of excess success for articles in *Psychological Science*. *Psychonomic Bulletin & Review*, 21(5), 1180-1187.
- Gelman, A., & Loken, E. (2013). *The garden of forking paths: Why multiple comparisons can be a problem, even when there is no "fishing expedition" or "p-hacking" and the research hypothesis was posited ahead of time*. Retrieved from http://www.stat.columbia.edu/~gelman/research/unpublished/p_hacking.pdf
- Green, P., & MacLeod, C. J. (2016). SIMR: an R package for power analysis of generalized

- linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493-498.
- Hertzog, M. A. (2008). Considerations in determining sample size for pilot studies. *Research in Nursing & Health*, 31(2), 180-191.
- Holtzman, W. H. (1950). The unbiased estimate of the population variance and standard deviation. *The American Journal of Psychology*, 63, 615-617.
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23(5), 524-532.
- Johnson, P. C., Barry, S. J., Ferguson, H. M., & Müller, P. (2015). Power analysis for generalized linear mixed models in ecology and evolution. *Methods in Ecology and Evolution*, 6(2), 133-142.
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 4, 863:1-12.
- Muthén, L. K., & Muthén, B. O. (2002). How to use a Monte Carlo study to decide on sample size and determine power. *Structural Equation Modeling*, 9(4), 599-620.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4), 867-872.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11), 1359-1366.

Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method.
Perception & Psychophysics, 33(2), 113-120.