A simple statistical framework for small sample studies

D. Samuel Schwarzkopf [1,2] & Zien Huang [1,2]

1. School of Optometry & Vision Science, University of Auckland, Grafton, Auckland, New Zealand
2. Experimental Psychology, University College London, London, United Kingdom

Abstract

Most studies in psychology, neuroscience, and life science research make inferences about how strong an effect is on average in the population. Yet, many research questions could instead be answered by testing for the universality of the phenomenon under investigation. By using reliable experimental designs that maximise both sensitivity and specificity of individual experiments, each participant or subject can be treated as an independent replication. This approach is common in certain subfields. To date, there is however no formal approach for calculating the evidential value of such small sample studies and to define a priori evidence thresholds that must be met to draw meaningful conclusions. Here we present such a framework, based on the ratio of binomial probabilities between a model assuming the universality of the phenomenon versus the null hypothesis that any incidence of the effect is sporadic. We demonstrate the benefits of this approach, which permits strong conclusions from samples as small as 2-5 participants and the flexibility of sequential testing. This approach will enable researchers to preregister experimental designs based on small samples and thus enhance the utility and credibility of such studies.

Keywords

statistics, inference, universality, replication, replicability, power, individuals, metascience

Introduction

Over the past decade, the fields of psychology and cognitive neuroscience have been embroiled in something of a methodological revolution, sparked by what is commonly referred to as the "replication crisis." Large-scale efforts revealed that a considerable proportion of published research findings fail replication attempts, the corner stone of scientific inference [1]. This has led to proposals for alternative statistical approaches, such as based on the estimation of effect sizes [2] or hypothesis tests based on Bayesian statistics [3–6]. It has also resulted in changes to the publication process, in particular the addition of peer review prior to data collection in the *Registered Report* format [7,8]. Vetting and adjusting the experimental design and analytical approach prior to even collecting the first data point can control for biases and researcher degrees of freedom arising from methodological flexibility [9]. It is also simply good practice to set *a priori* thresholds for the evidence required to reach a conclusion.

These changes have thus far focussed primarily on the statistical framework of estimating population tendencies from limited samples. Such inferences are necessary for answering many research questions, such as quantifying the effect of a medical treatment or a psychological manipulation. However, this often goes together with requiring ever larger sample sizes to test increasingly more subtle effects. When used unwisely, this approach is wasteful and could prove unsustainable. There are only limited resources for scientific research; expecting vast samples (possibly collected by multi-site consortia) necessitates diverting funds that could otherwise have been shared by several teams to address a range of questions. The inevitable consequence is that scientific progress will slow down. However, contrary to the ingrained belief of many journal editors and reviewers, group studies based on large sample sizes are not necessary for all kinds of research [10]. Indeed, for some research the mindless use of group statistics [11] can even hinder the interpretation of results [12].

The purpose of psychology and neuroscience research is to discover the mechanisms of the mind and brain. To answer many research questions, the average population effect is at best uninformative, or even irrelevant. Sensory psychophysics studies regularly employ a small sample (small-N) approach in which hypotheses are tested within each individual observer. This considers each observer as an independent replication. Some studies can demonstrate convincing results with a sample as small as two participants. Rather than inferring the average effect in the population, these studies ask whether the phenomenon in question is universal to human perception. Provided the individual tests are reliable, demonstrating their existence in a small number of people is ample evidence. Importantly, while this approach is common in psychophysical studies, it should also be suitable in many other scenarios, such as neuroimaging or electrophysiological studies, and beyond.

But how many participants are necessary to confirm the existence of a phenomenon? And how does this depend on the reliability of individual tests? Without concrete answers to these questions, small sample studies will continue to run afoul of ill-advised editorial decisions and reviewer recommendations. What is worse, the current climate fosters widespread misconceptions about what makes good science.

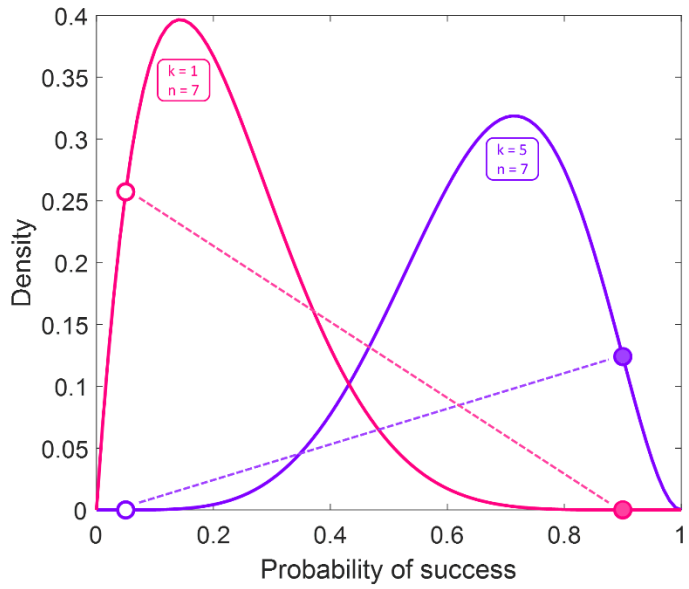

Figure 1. Binomial probability density functions for achieving k=1 success (pink) or k=5 successes (purple) out of n=7 total tests (subjects) plotted against the success probability. Filled discs denote the binomial probability when success probability is 0.9, a common assumption for sensitivity. This is the likelihood of the alternative hypothesis, $P(D|H_1)$. Open discs denote the binomial probability when assuming successful tests occur by chance, the false positive rate 0.05 of the individual tests. This is the likelihood of the null hypothesis, $P(D|H_0)$. The evidence ratio, ε, is the ratio of these likelihoods. The ratio is large/small when the alternative hypothesis is more/less likely than the null hypothesis (visualised by dashed lines).

Binomial model comparison

Here we present a basic framework for estimating the statistical evidence about results from small-N designs. It compares the likelihood of two opposing models: the alternative (experimental) hypothesis, H_1, assumes that the phenomenon of interest is a universal feature across all subjects. Any failure to detect it in a subset of individual tests should merely be due to lacking sensitivity of these tests, perhaps because of lapses in performance or other individual differences skewing the measurement. In contrast, the null hypothesis, H_0, assumes that the phenomenon is not universal, and that any individuals showing a significant effect are either rare outliers or random flukes.

The model comparison is then conducted via a likelihood ratio, similar to what is commonly done for statistical inference via Bayes Factors or likelihood tests [4,5,13]. To determine the likelihood of H_1 we calculate $P(D|H_1)$, the binomial probability density of obtaining the observed number of significant subjects, *k*, out of *n* participants in total, given the sensitivity of the individual test. The sensitivity is 1-β, where β is the false negative rate. The likelihood of H_0 is $P(D|H_0)$, the binomial probability of observing the same result (*k* successful out of *n* tests in total) under the assumption that the effect occurred by chance. The chance probability of the individual tests is their significance threshold, or the false positive rate α. To illustrate, Figure 1 plots two example binomial probability density functions for a set of 7 participants out of whom either only 1 or 5 yielded significant findings. The evidence ratio, ε, is then given by,

$$\varepsilon = \frac{P(D|H_1)}{P(D|H_0)} = \frac{P(k|n, 1-\beta)}{P(k|n, \alpha)}$$

This quantifies how much more likely the observed result D is under H_1 than under H_0. In pragmatic terms, if ε is high we would expect the result to replicate when running another subject. Since the calculation of the evidence relies on the error rates, α and β, this necessitates an experimental design that can minimize these errors. Determining the false positive rate is relatively straightforward,

especially in designs that use standard significance tests at the level of individual subjects, such as through a p-value or a confidence interval. Deriving the false negative rate is more complicated, especially as sensitivity can vary between subjects.

To illustrate, let us assume perfect sensitivity. If the phenomenon is universal, the test will always detect it. If the test also had perfect specificity, $\alpha=0$, we would only need to test a single subject to draw a conclusion about whether the phenomenon exists. Such a perfect test is rare, if it exists at all; even highly reliable procedures will carry at least a minor risk of false positives. Thus, assuming a conventional significance threshold of $\alpha=0.05$ for the individual tests, after running a single subject we can obtain the following evidence ratios: if this single test fails to reach significance, this rules out the existence of the effect because with perfect sensitivity we should have confirmed H_1. However, if the individual test is significant, this indicates that H_1 is 20 times more likely than H_0. Such an evidence ratio is relatively strong, although maybe not sufficiently compelling to stop the experiment after only a single subject.

What about running two subjects? If both test significant, the evidence is $\varepsilon=400$ ($=0.05^{-2}$), which seems like a convincing result. But because the sensitivity of the individual test is perfect, any other result would support H_0, however, and completely rule out that the phenomenon is universal. Naturally, perfect sensitivity is just as improbable as perfect specificity. Even the most reliable test sometimes fails to detect a phenomenon. Statistically, this means we might assume a sensitivity somewhat less than perfect, perhaps 0.95. Now, if we test only two subjects and find a significant effect in both, this produces an evidence ratio of $\varepsilon=361$. Conversely, finding no significant effects, does not rule out its existence but gives $\varepsilon=1/361$, approximately 0.0028. These are both still very conclusive results for H_1 or H_0, respectively. Crucially, finding only one of the two subjects was significant means that the result is completely inconclusive ($\varepsilon=1$). In this case, more subjects are required to make an inference. If we tested a third subject, and their result was also significant, this would again produce a moderate $\varepsilon=19$. Three significant out of four subjects would bring back compelling evidence at $\varepsilon=361$.

What this toy example demonstrates is that the small-N approach naturally lends itself to sequential testing. This is unsurprising, given that the logic is essentially the same as for inference using Bayes Factors, which is commonly used for sequential sampling in group studies [14]. Crucially, this means that researchers can (and should) define a threshold for the evidence ratio needed to make an inference. The choice of threshold is up to researchers and depends on the nature of the study. Confirming a well-known sensory aftereffect is probably compelling with $\varepsilon\geq16$ (i.e., corresponding to a conventional significance threshold of 0.05 and sensitivity of 0.8). In contrast, testing extraordinary claims of precognitive abilities, as in a study that was partly responsible for triggering the replication crisis in psychology [15], should probably demand more extraordinary evidence. As a rule of thumb, we chose a relatively conservative threshold of 50 (or 1/50=0.02, respectively).

Determining the sensitivity

One problem common to experimental design, not only for this small-N approach but also classical group studies, is how to determine the sensitivity, or the statistical power, of the experiment. Calculating power requires a predefined effect size of interest. This is often based on the effect sizes reported in previous research, although these may not always seem appropriate because previous studies used different experimental paradigms or different dependent variables. Most previous research is also affected by publication bias, irrespective of the methodological rigour or controls put in place [16–19]. There may be an unknown number of attempts at similar experiments whose results were not strong enough to make it into the literature. This means that predicted effect sizes based on previous research are likely inflated. A more principled approach could be to determine an effect size of interest based on the practical significance of an effect.

The latter type of consideration is particularly crucial for the small-N approach. In most contexts, it should be possible to design the tests in individual subjects such that they are highly sensitive to reveal practically meaningful effects. For instance, when estimating an illusion magnitude (subjective perceptual bias) by fitting a psychometric curve, a practical consideration could be that the detectable bias should exceed the just noticeable difference (objective discrimination ability). The experiment should therefore collect enough trials to ensure that such a bias can be detected reliably. Similar considerations could determine the sensitivity of neuroimaging studies, e.g., the correlation between voxel response patterns from a brain region within a single subject. Power analysis can determine the minimum correlation detectable with the number of voxels in a brain region. Critically, because sensitivity varies between subjects, this estimate should aim to maximize sensitivity for the weakest link, that is the subject thought to have the lowest sensitivity, due to the smallest number of voxels or trials, etc.

However, even in the absence of a clear estimate of individual sensitivity, it is possible to estimate the evidence ratio by plotting the full range of possible sensitivities, i.e., between the false positive rate α and 1 (we assume that in well-designed studies, sensitivity should always exceed α). Evidence curves for a range of typical sample sizes (Figure 2) show that when the evidence favours the null hypothesis, that is, when only a few subjects test significant, the support for H_0 increases steadily with sensitivity. In contrast, when most subjects in a sample test significant, the evidence peaks at an intermediate sensitivity and then declines. As already discussed, unless all subjects test significant, the evidence ratio at sensitivity of 1 must be 0. Therefore, highly sensitive tests also imply a high likelihood that all subjects test significant and any divergence from this result provides evidence against H_1. The peak of the curve in fact corresponds to the observed proportion of significant findings: e.g., with 2 significant out of 3 subjects the evidence curve peaks at ~0.667 sensitivity.

Inspection of these evidence curves allows an insight into the overall evidence for H_1, even when the sensitivity is uncertain. By determining the range of likely sensitivities of the individual tests one can identify the strength of the evidence. If a researcher is confident that their individual tests have near-perfect sensitivity, then any result but an overwhelming majority of subjects testing significant argues against the universality of the phenomenon.

For most experiments, such extreme sensitivity is however unachievable; therefore, a fraction of non-significant subjects can still provide strong evidence for H_1. With a sample size of n=2 (Figure 2A), if both subjects test significant this would pass the $\varepsilon>50$ criterion at a relatively low sensitivity of only ~0.354. However, if no subject tests significant, this would not reach the evidence threshold for H_0 of $\varepsilon<0.02$, unless sensitivity was at least ~0.866. Clearly, if the researcher has high confidence in the individual tests, such a result convincingly argues against the universality of the phenomenon. Finally, if only one of the two subjects tests significant, the evidence curve always stays between 0.02 and 50. More data must therefore be collected regardless of the sensitivity of individual tests.

The relationship is similar for larger sample sizes (Figure 2B-F). Larger samples simply have the consequence of shifting down the required sensitivity to reach conclusive evidence for either hypothesis. Just as is the case for classical group statistics, it is easier to obtain evidence supporting the existence of an effect than it is to support its absence. At n=10, the necessary sensitivity for obtaining $\varepsilon>50$ if all subjects test significant is ~0.074. This is an overwhelming false negative error rate and suggests the need for a better experimental design. Conversely, if none of the ten subjects test significant, the sensitivity needed to obtain support for H_0 at $\varepsilon<0.02$ is ~0.357, still a relatively insensitive test. This example underlines the fact that larger samples are not superior in this approach; rather they only benefit poorly sensitive studies.

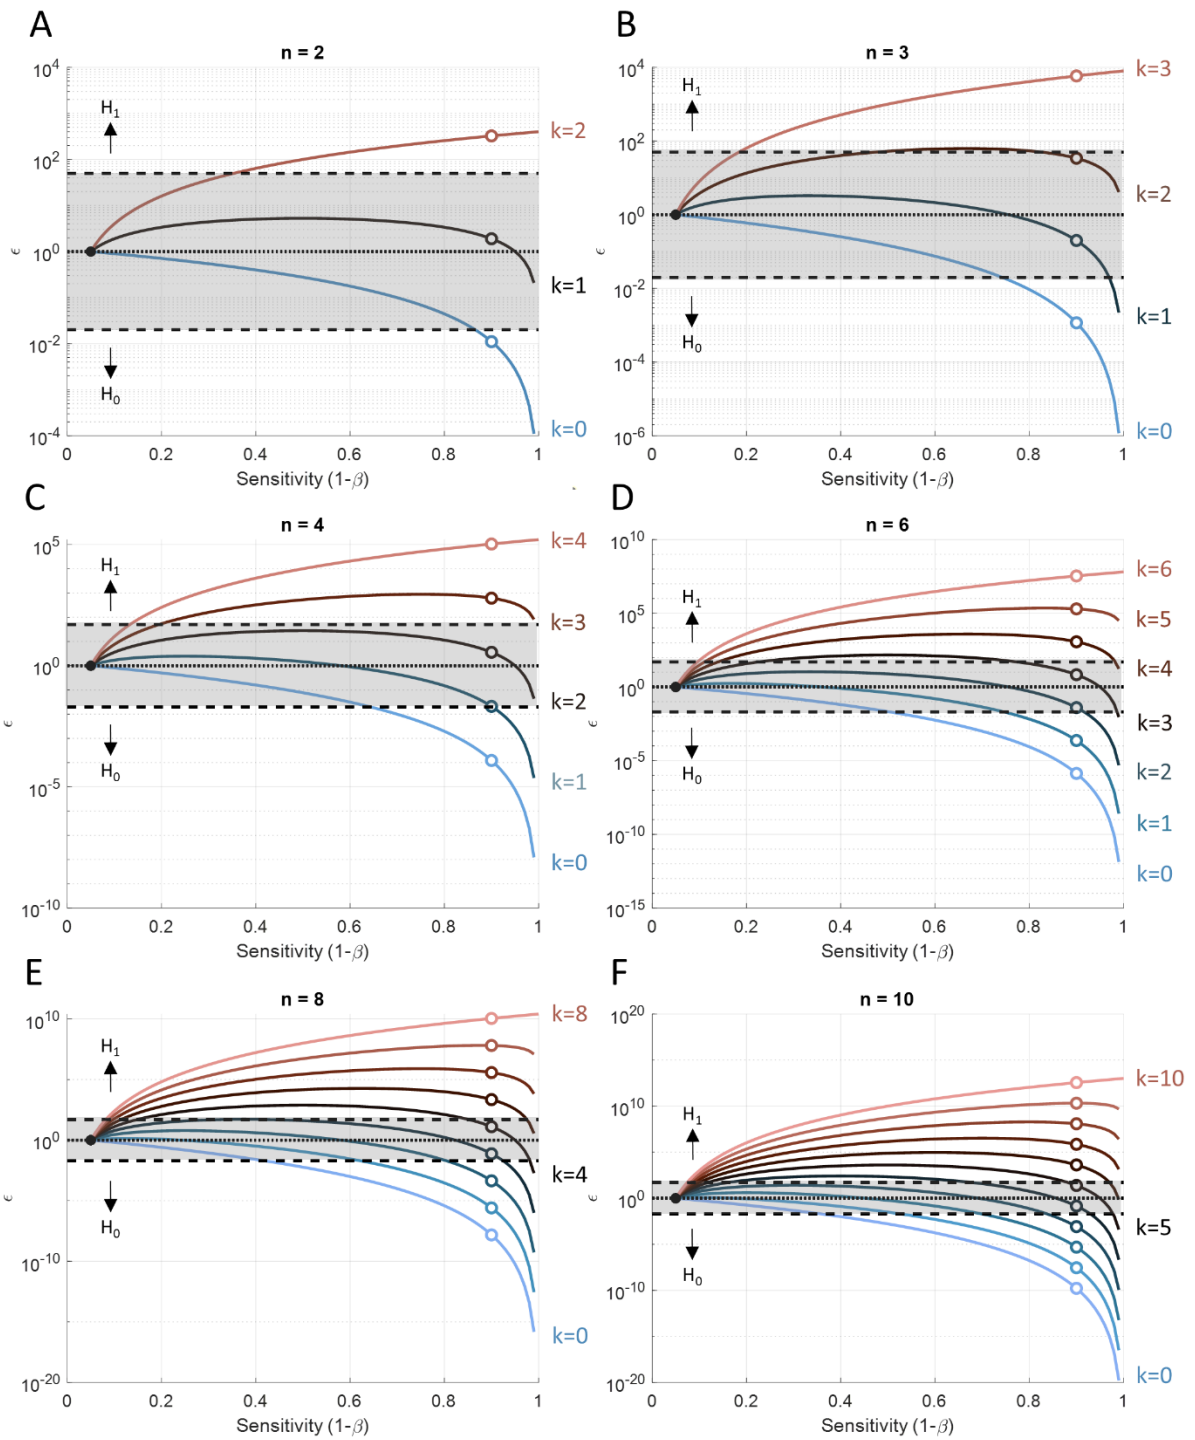

Figure 2. Statistical evidence plotted as a function of sensitivity of individual tests. Each curve denotes the evidence ratio ε obtained for a given number of significant tests in the sample, k, ranging from no significant tests (k=0, blue) to all tests significant (pink). Curves are shown for sample size n=2 (A), n=3 (B), n=4 (C), n=6 (D), n=8 (E), and n=10 (F). The dashed black lines and grey shaded areas indicate the region of inconclusive evidence where 0.02<ε<50. Evidence ratios above or below this region indicate compelling support for H_1 or H_0, respectively. For reference, open circle symbols indicate sensitivity of 0.9.

We summarise this relationship in Figure 3. This shows across a range of sample sizes the sensitivity required for passing the evidence thresholds for H_1 and H_0, respectively, when either all or none of the subjects are significant. As the name implies, the small-N approach is designed to provide conclusive evidence for small samples. It demands maximising the sensitivity of individual tests. If high sensitivity can be achieved, a conclusive study can be conducted with substantially fewer than ten subjects. However, if sensitivity cannot be optimised, the approach is ill-advised, and a conventional group average design is more suitable.

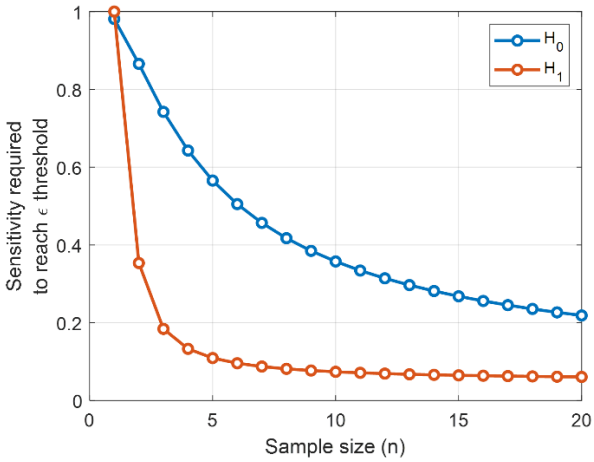

Figure 3. Sensitivity necessary for obtaining conclusive evidence with a given sample size. Curves plot the required sensitivity needed for $\varepsilon>50$ for supporting H_1 when all subjects test significant (red) or $\varepsilon<0.02$ for supporting H_0 when no subject tests significant (blue).

Relationship to group statistics

By now, it has probably become apparent that the small-N approach is particularly suitable for robust effects. If the phenomenon manifests in most of the sample, this strongly supports its existence. Naturally, a highly consistent effect across the whole sample is also likely to be significant in a classical group test. Why then should we use the small-N approach instead of simply using classical group statistics on the small sample? One reason is that many investigators regard significant findings from small samples with suspicion, even if the results are compelling. More importantly, as we discussed, the small-N approach asks a different question than group statistics. The small-N approach is indicated whenever we make inferences about universal features common to all subjects rather than quantifying the magnitude of this feature. To conclude that typically-developed humans are born with an appendix, we need only confirm they have this vestigial organ; we would not test that the average human appendix is longer than 0 cm.

To illustrate the superiority of the small-N approach for such research questions, we sought to replicate a well-established effect from perceptual psychophysics, known as the oblique effect [20–22]: observers' ability to discriminate the orientation of oblique orientations (that is, near 45° or 135°) is considerably worse than performance near cardinal orientations (0° and 90°, corresponding to horizontal and vertical). This difference in performance is thought to reflect differences in the neural populations in visual cortex coding for the different orientations [21,23], possibly related to the statistical distribution of oriented edges in the natural environment [24]. We are inclined to treat this is a fact about human visual performance.

We collected orientation discrimination data for the oblique and vertical orientations in a sample of n=17 observers (see Supplementary Information). We quantified this using a psychometric curve fit where the bandwidth is an indication of discrimination ability: narrower bandwidths correspond to better performance. This confirmed our prediction that average bandwidth for judging orientations

near 90° was smaller than near 45° obliques (one-tailed t-test for $\mu>0$, $t_{16}=13.88$, $p<0.0001$; Bayes factor with a default Cauchy prior scaled with $r=\sqrt{2}/2$, BF_{10}= 74,700,000), an example of a robust effect. Conversely, comparison of the two oblique orientations revealed no significant difference in performance ($t_{16}=-1.14$, $p=0.8644$) and Bayesian inference favours the null hypothesis that bandwidths near 45° were not greater than near 135° ($BF_{10}=0.13$). We use this dataset as an example of a compelling null result.

What would have happened had we analysed this experiment using the small-N approach? Our psychometric curve fit also allowed us to make statistical inferences for each individual observer. Doing so revealed that the oblique effect (bandwidths near 45°>90°) was significant at $p<0.05$ in all but one observer. Even this non-significant observer replicated the effect numerically. Under the assumption that this phenomenon is a replicable effect that should manifest in all observers, we take this as an estimate of the sensitivity of our test, 0.941. In contrast, the null effect (bandwidths near 45°>135°) was only significant in one observer. This corresponds to ~0.059, close to the nominal false positive rate expected of the test.

Thus, using $\alpha=0.05$ and a sensitivity of $1-\beta=0.941$, we calculated the evidence ratio when incrementing the sample size sequentially in the order individuals were tested (Figure 4A, purple curve). This showed that even with two observers, the evidence for the oblique effect already passed the threshold of $\varepsilon>50$. Testing more observers only further increased the evidence. By the time we reached observer 12, whose results failed to reach significance, this could hardly dent the evidence. Importantly, for the null effect we also reached conclusive evidence in favour of H_0 with two subjects already (Figure 4B, purple curve). By the time, we ran observer 15, who tested significant despite expecting no difference, the evidence for the null hypothesis was already overwhelming. Essentially, using the small-N approach we could have ended data collection after two observers already – unheard of in conventional group studies.

In comparison, incrementing sample sizes using classical group statistics would have forced us to continue the experiment (Figure 4, orange curves). We only report Bayes Factors for this analysis because sequential testing skews frequentist statistics [9,14], but see also [25]. For the oblique effect (Figure 4A), after 3 observers we would have obtained a Bayes Factor above the criterion of $BF_{10}>6$ used as evidence threshold in the *Registered Reports* format by many journals (e.g., *Cortex:* https://rr.peercommunityin.org/about/pci_rr_friendly_journals#h_49206884940316184193307 27). This would therefore have entailed testing one more observer. This may not seem like a big difference; however, even with a strong Bayes Factor like this ($BF_{10}=14.1$), experience tells us that few people would find this result compelling. In fact, we wager that this would lead to immediate desk rejection by the handling editor at some journals. On the other hand, the small-N approach demonstrates that even two observers testing significant is already extremely convincing evidence.

However, the true benefit of the small-N approach becomes clear for the null effect. The group analysis only produced any evidence in favour of the null hypothesis ($BF_{10}<1/6$) by the time we had tested observer 11. We would have needed to collect more than five times the sample size required for reaching the same conclusion using the small-N approach. The Bayes factors for the null effect were also highly unstable, jumping up and down as the sample size increased. After testing observer 15, the one whose results showed an unexpected significant difference at the individual level, the Bayes factor came dangerously close to falling above the criterion for H_0 ($BF_{10}=0.164$). With small samples, individual effect sizes can have a profound effect on the sample mean, resulting in considerable fluctuations in the evidence of group tests. In the small-N approach, however, the evidence for H_0 steadily increased, with observer 15 causing only a tiny blip. Note also that these conclusions do not strongly depend on the chosen evidence threshold. The situation would be similar for a criterion of $BF_{10}>3$ or $BF_{10}<1/3$, respectively, especially considering that our criterion for the small-N approach of $\varepsilon>50$ or $\varepsilon<0.02$ could be considered highly conservative.

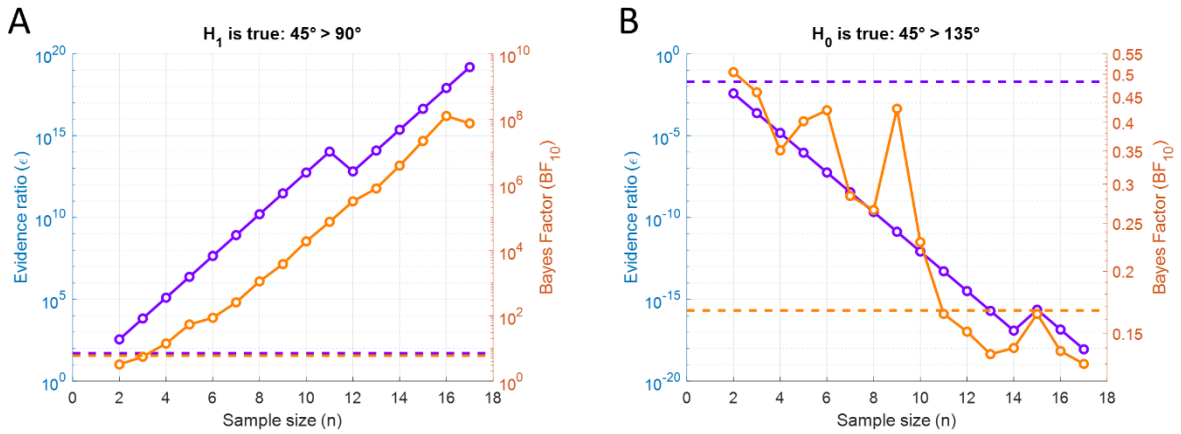

Figure 4. Statistical evidence in oblique effect experiment as a function of sample size. Evidence ratio ε (purple) of the small-N test, and Bayes Factor BF_{10} of a one-sample t-test for μ>0 (orange), were plotted against the sequentially increasing sample size, shown separately for the assumed true effect (A) and the assumed null effect (B). Note separate y-axes for the two measures. Dashed horizontal lines denote the thresholds for conclusive evidence supporting H_1 (ε>50, BF_{10}>6) or H_0 (ε<1/50, BF_{10}<1/6), respectively.

One caveat is that the small-N approach depends on the sensitivity of individual tests. In our example, we estimated sensitivity from the assumed false negative rate. We could only do so because he had collected a larger sample, and because we were already convinced that the oblique effect is real. In typical experiments, however, one would need to rely on other estimates. For a psychophysical study such as this, sensitivity is usually high, provided enough trials are collected per observer. Through simulations one could determine the required number of trials for estimating the perceptual effect, although this can be a relatively complex because psychophysical performance depends on various factors. However, as already explained, even without precise knowledge of the achieved sensitivity, inspection of the evidence curves (Figure 2A) permits the conclusion that two out of two observers showing a significant oblique effect strongly support its existence; the evidence is already conclusive even with comparably low sensitivity (Figure 3, red curve). For supporting the null effect, much higher sensitivity would be required (Figure 3, blue curve), so we can only draw this conclusion with highly sensitive tests. To support a null result, one would likely need to collect a few more observers, around 4 or 5 in total. However, the required number would still be considerably smaller than the 11 we needed using classical group statistics.

Examples from the literature

We suspect most researchers will consider samples of only two subjects to be very small. However, there are examples of this in the literature. A seminal psychophysical study [26] used only two observers – who also happened to be authors of the article – to characterise the parameters that govern performance in the contour integration ('pathfinder') task where subjects must detect the presence of a group of aligned Gabor patches in a cluttered field of randomly oriented distractors. Despite the small sample of non-naïve observers, the results are clear and compelling. It would be difficult to estimate the sensitivity of these tests, but we can be confident that it far exceeds what is necessary to detect the reported effects. The study reported five experiments, four of which show the same significant effect in both observers and the fifth shows a convincing null result.

There are also good examples from the neuroimaging literature. Another famous study used an encoding model to demonstrate that the identity of natural images can be accurately predicted based

on response patterns from early visual cortex [27]. They demonstrated this result in two subjects. Nobody would doubt this result, given the convincing reliability of these effects within each subject. While this may be an extreme example, the same concept applies to other neuroscience findings whenever many data points are used to make an inference at the level of individual subjects, for example when comparing the response of a set of independently localised voxels, a population of neurons in single-cell recordings, or a representational similarity analysis using patterns of responses.

A corollary is that methods like representational similarity analysis [28] are more powerful than frequently used multivariate classification decoders [29–31]. The latter typically predict condition labels (stimulus categories, mental states) based on patterns of brain activity on a trial-by-trial basis. The sensitivity of this analysis therefore depends solely on the number of trials producing the classification accuracy for each subject. This number is typically small, and so the sensitivity is also low. In contrast, in a representational similarity analysis we correlate the voxel response patterns of independent stimulus conditions. The sensitivity is given by power analysis for correlations with the number of voxels in the pattern, based on classical power calculations. Using the small-N approach, a significant correlation in 2-3 subjects could already be compelling evidence, whereas classification accuracy in a typical design is unlikely to be significantly greater than chance in all three subjects.

Finally, the small-N approach can also provide meaningful answers where the outcome of conventional group statistics can be misleading. Claims of precognitive abilities in classical psychology experiments were supported by significant group statistics [15]. Specifically, experiment 1 of this study reported performance for guessing the outcome of future trials with 53.1%, significantly above the chance level of 50% (t_{99}=2.51, p=0.01). The experiment used 100 subjects, each of whom performed 18 or 12 trials, respectively, in the critical condition. The average precognitive ability therefore constitutes less than a single trial above chance. The author's argument was that to detect a phenomenon as subtle as this requires averaging performance over many subjects.

However, by instead increasing the number of trials performed by each subject, the individual sensitivity should be sufficient to detect extremely subtle effects, even smaller than the reported finding. Thus, the same experiment could have been conducted using the small-N approach. Admittedly, this would entail running thousands of trials in each person. But the evidence of finding significant above-chance performance in only 3-4 subjects with a high-powered procedure would be far harder for sceptics to ignore than a significant group effect in a sample of 100 subjects. On this note, a recent replication attempt of this experiment used a methodologically highly rigorous procedure and found no evidence for precognition [32]. This was based on a sample of n=2220 subjects. It could be more efficient to test four people 200 times, than running an experiment on such a huge sample; but of course, there could be reasons why only a small number of trials can be collected per subject; for instance, some experiments are specifically designed with one critical trial in mind [33].

Conclusions

We introduced a simple framework for estimating the statistical evidence of studies using small samples, typically comprising fewer than ten subjects. Despite justified concerns about the statistical power and replicability of research in the psychology, neuroscience, and other areas of life sciences, obviously not all studies in those fields require large sample sizes. Any experiment that can provide reliable data from individual participants or subjects, based on large numbers of trials or observations, is suitable for small-N analysis. Rather than estimating the average population effect, the small-N approach treats each subject as an independent replication. Most importantly, it does not require extensive knowledge of complex statistical methods because it is simply based on the binomial probability function. This can be calculated using most standard statistical analysis packages or even online calculators. We provide a set of MATLAB functions with this article for the readers' convenience.

Small sample studies are commonplace in several subfields of neuroscience research, ranging from psychophysical studies to neurophysiological recordings in animal models. While this has been pointed out previously [10], to date no straightforward frameworks for interpreting the evidential value of such studies have been described. By formalising this approach, we hope to expand the use of small sample approaches to other research contexts where the benefits of this approach have thus far not been realised. Moreover, we hope to raise awareness among researchers (including journal editors and reviewers) that such an approach can yield robust and meaningful results.

Our example, using the well-established oblique effect [20,21], shows that the small-N approach produced conclusive results with only two subjects, while conventional group analysis would have required more than five times this number to support an assume null result. Consider the waste of resources, work hours, and funds that this might entail in a real experiment. Human volunteers are precious. Running unnecessary experiments is counterproductive: many participants who might be disinclined to return for more experiments could have participated in other studies. It is also unethical, and this issue becomes even greater in animal studies, where researchers have an obligation to reduce the number of subjects and minimise their suffering.

We note that a similar approach to ours has recently been described which uses a Bayesian framework to estimate the prevalence of a phenomenon in the population [12,34]. This approach complements our small-N approach but there are a few points of divergence. First, as the name implies the prevalence approach quantifies the frequency of the phenomenon, and this could be relatively low. The prevalence approach could therefore be useful to revealing individual differences or significant outliers [12]. In contrast, our small-N approach assumes that an effect is universal (under the alternative hypothesis), or that it does not exist and any incidence of it either arose by chance or because it is rare (null hypothesis). A corollary is that the prevalence approach may still require relatively larger samples to draw confident conclusions about prevalence. Conversely, the small-N framework is literally designed to minimise the number of required subjects.

Depending on the situation, the small-N approach also allows us to gain insights into scenarios with individual differences. For example, our empirical demonstration was based on the rationale that the oblique effect manifests in all participants. We found a single observer without a significant result. This is most likely a false negative; while sensitivity was presumably high, there is a small chance that one test will fail to reach significance. The evidence ratio shows this result is far more likely assuming the effect's universality than under the null hypothesis of random false positives. However, it is also possible that the effect is common but *not* universal. Repeating the experiment on this observer could put this hypothesis to the test: if the effect is universal, it is highly likely significant on the second test. However, if the effect were due to true individual differences, it would be more likely that it again fails to reach significance.

Purists in statistical philosophy might balk at our mixture of frequentist statistical tests at the first (individual subject) level and a likelihood approach at the second level aggregating the evidence. However, we disagree with this objection on pragmatic terms. For one thing, our approach is extremely versatile. While most examples used here involve frequentist statistics at the first level, this is not required. Our approach simply quantifies the number of subjects/tests successfully confirming a hypothesis. This can take many forms that do not involve formal significance tests; for example, a success could simply be based a performance criterion or some other numerical threshold. This criterion should of course be chosen to maximise sensitivity and minimise false positives. That said, in many situations one could construct a pure likelihood approach to reach similar conclusions as we do with our approach. For example, one study used a Bayesian hierarchical model to quantify the evidence of highly consistent findings like the Stroop effect [35]. More generally, instead of individual significance tests one could use likelihood tests at the first level. Quantifying the evidence would then simply involve multiplying the individual likelihoods. This approach is more sensitive because it does not binarise the individual tests. However, it necessitates that a likelihood ratio can be easily calculated

for individual tests. In contrast, the small-N approach we describe here is easily understood without extensive knowledge of statistics and thus suitable to a wide user base.

Importantly, the small-N approach is not a replacement for conventional group statistics. Rather it adds to the toolbox of scientific research when the research question is suitable. Many research questions could have been better answered using this approach than conventional statistics. However, in the same vein, there are numerous questions for which there is no alternative but to use group statistics, whatever the flavour is conventional null-hypothesis significance testing, Bayesian hypothesis testing [3–6], or effect size estimation [2]. If the goal of the research is to have a precise knowledge of the magnitude of the effect, the latter is probably the best option. Often this will necessitate very large samples [36]. Similarly, even if the exact population effect size is not crucial, but an estimate is required to interpret the findings, then conventional group statistics are necessary. Naturally, if individual subjects only contribute singular measures (such as a questionnaire score, a behavioural performance level, or a neurophysiological amplitude, etc.) only the conventional group analysis is possible. The choice of statistical approach therefore depends on the research question and the experimental design. We hope that our exposition of the small-N approach will encourage researchers to consider how they could apply it to answer their research questions instead of mindlessly relying on conventional group statistics.

It should also go without saying that the small-N approach is not an excuse for lacking methodological rigour. In writing this article we in fact aim to enhance the rigour of studies using small samples. For the first time, this framework allows researchers to set predefined levels of evidence necessary to make inferences from their small-N data. This could also be useful for replication efforts. Because the small-N approach critically depends on the sensitivity of individual tests, it focusses researchers on maximising sensitivity. The approach also directs attention to other aspects of the experimental process, such as clearly defining exclusion criteria or pre-processing parameters. Despite their advantages, small-N studies are obviously not immune from the adverse effects of flexible methods. Experimental procedures and analysis methods should normally be determined – and ideally preregistered – prior to data collection. Thus far, without any way to determine evidence thresholds small sample studies have been automatically precluded from enjoying methodological advances like *Registered Reports* in many journals [7,8]. Our framework now offers researchers greater control of biases arising from methodological flexibility and boost the credibility of scientific research based on small samples.

Data availability

Analysis code and data from the example experiment are publicly available at https://osf.io/bmzc3/

Acknowledgements

We thank Zoltan Dienes, Benjamin de Haas, and Will Harrison for useful comments on this manuscript. We also thank the editors and reviewers at multiple journals who inspired us to write this article by insisting that only ever-larger samples can yield meaningful results.

References

1. Open Science Collaboration. 2015 PSYCHOLOGY. Estimating the reproducibility of psychological science. *Science* **349**, aac4716. (doi:10.1126/science.aac4716)

2. Cumming G. 2014 The new statistics: why and how. *Psychol. Sci.* **25**, 7–29. (doi:10.1177/0956797613504966)

3. Dienes Z. 2014 Using Bayes to get the most out of non-significant results. *Quant. Psychol. Meas.* **5**, 781. (doi:10.3389/fpsyg.2014.00781)

4. Rouder JN, Speckman PL, Sun D, Morey RD, Iverson G. 2009 Bayesian t tests for accepting and rejecting the null hypothesis. *Psychon. Bull. Rev.* **16**, 225–237. (doi:10.3758/PBR.16.2.225)

5. Wagenmakers E-J. 2007 A practical solution to the pervasive problems of p values. *Psychon. Bull. Rev.* **14**, 779–804.

6. Wagenmakers E-J *et al.* 2018 Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychon. Bull. Rev.* **25**, 35–57. (doi:10.3758/s13423-017-1343-3)

7. Chambers CD. 2013 Registered reports: a new publishing initiative at Cortex. *Cortex* **49**, 609–610. (doi:10.1016/j.cortex.2012.12.016)

8. Chambers CD, Dienes Z, McIntosh RD, Rotshtein P, Willmes K. 2015 Registered reports: realigning incentives in scientific publishing. *Cortex* **66**, A1-2. (doi:10.1016/j.cortex.2015.03.022)

9. Simmons JP, Nelson LD, Simonsohn U. 2011 False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychol. Sci.* **22**, 1359–1366. (doi:10.1177/0956797611417632)

10. Smith PL, Little DR. 2018 Small is beautiful: In defense of the small-N design. *Psychon. Bull. Rev.* **25**, 2083–2101. (doi:10.3758/s13423-018-1451-8)

11. Gigerenzer G. 2004 Mindless statistics. *J. Socio-Econ.* **33**, 587–606. (doi:10.1016/j.socec.2004.09.033)

12. Ince RAA, Kay JW, Schyns PG. 2022 Within-participant statistics for cognitive science. *Trends Cogn. Sci.* **26**, 626–630. (doi:10.1016/j.tics.2022.05.008)

13. Jeffreys H. 1961 *Theory of probability.* Oxford, UK: Oxford University Press.

14. Rouder JN. 2014 Optional stopping: No problem for Bayesians. *Psychon. Bull. Rev.* (doi:10.3758/s13423-014-0595-4)

15. Bem DJ. 2011 Feeling the future: experimental evidence for anomalous retroactive influences on cognition and affect. *J. Pers. Soc. Psychol.* **100**, 407–425. (doi:10.1037/a0021524)

16. Kühberger A, Fritz A, Scherndl T. 2014 Publication Bias in Psychology: A Diagnosis Based on the Correlation between Effect Size and Sample Size. *PloS One* **9**, e105825. (doi:10.1371/journal.pone.0105825)

17. Masicampo EJ, Lalande DR. 2012 A peculiar prevalence of p values just below .05. *Q. J. Exp. Psychol. 2006* **65**, 2271–2279. (doi:10.1080/17470218.2012.711335)

18. Newcombe RG. 1987 Towards a reduction in publication bias. *Br. Med. J. Clin. Res. Ed* **295**, 656–659.

19. Rosenthal R. 1979 The file drawer problem and tolerance for null results. *Psychol. Bull.* **86**, 638–641.

20. Appelle S. 1972 Perception and discrimination as a function of stimulus orientation: the 'oblique effect' in man and animals. *Psychol. Bull.* **78**, 266–78.

21. Furmanski CS, Engel SA. 2000 An oblique effect in human primary visual cortex. *Nat. Neurosci.* **3**, 535–6.

22. Vogels R, Orban GA. 1985 The effect of practice on the oblique effect in line orientation judgments. *Vision Res.* **25**, 1679–87.

23. Coppola DM, White LE, Fitzpatrick D, Purves D. 1998 Unequal representation of cardinal and oblique contours in ferret visual cortex. *Proc. Natl. Acad. Sci.* **95**, 2621–2623. (doi:10.1073/pnas.95.5.2621)

24. Coppola DM, Purves HR, McCoy AN, Purves D. 1998 The distribution of oriented contours in the real world. *Proc. Natl. Acad. Sci.* **95**, 4002–4006. (doi:10.1073/pnas.95.7.4002)

25. Lakens D. 2014 Performing high-powered studies efficiently with sequential analyses. *Eur. J. Soc. Psychol.* **44**, 701–710. (doi:10.1002/ejsp.2023)

26. Field DJ, Hayes A, Hess RF. 1993 Contour integration by the human visual system: evidence for a local 'association field'. *Vision Res.* **33**, 173.

27. Kay KN, Naselaris T, Prenger RJ, Gallant JL. 2008 Identifying natural images from human brain activity. *Nature* **452**, 352–355. (doi:10.1038/nature06713)

28. Kriegeskorte N, Mur M, Bandettini P. 2008 Representational similarity analysis - connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* **2**, 4. (doi:10.3389/neuro.06.004.2008)

29. Cox DD, Savoy RL. 2003 Functional magnetic resonance imaging (fMRI) 'brain reading': detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage* **19**, 261–70.

30. Haynes J-D, Rees G. 2005 Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat. Neurosci.* **8**, 686–91. (doi:10.1038/nn1445)

31. Kamitani Y, Tong F. 2005 Decoding the visual and subjective contents of the human brain. *Nat. Neurosci.* **8**, 679–85. (doi:10.1038/nn1444)

32. Kekecs Z *et al.* 2023 Raising the value of research studies in psychological science by increasing the credibility of research reports: the transparent Psi project. *R. Soc. Open Sci.* **10**, 191375. (doi:10.1098/rsos.191375)

33. Dijkstra N, Fleming SM. 2023 Subjective signal strength distinguishes reality from imagination. *Nat. Commun.* **14**, 1627. (doi:10.1038/s41467-023-37322-1)

34. Ince RA, Paton AT, Kay JW, Schyns PG. 2021 Bayesian inference of population prevalence. *eLife* **10**, e62461. (doi:10.7554/eLife.62461)

35. Rouder JN, Haaf JM. 2021 Are There Reliable Qualitative Individual Difference in Cognition? *J. Cogn.* **4**, 46. (doi:10.5334/joc.131)

36. de Winter JCF, Gosling SD, Potter J. 2016 Comparing the Pearson and Spearman correlation coefficients across distributions and sample sizes: A tutorial using simulations and empirical data. *Psychol. Methods* **21**, 273–290. (doi:10.1037/met0000079)

Supplementary Material

Methods

Participants

Twenty participants (age range: 18-48 years, 15 female) with normal or corrected-to-normal visual acuity participated in the study. All gave written informed consent. The University of Auckland Human Participants Ethics Committee approved experimental procedures. Three participants were excluded from analysis because they failed to comply with task instructions.

Stimuli

Stimuli were generated using MATLAB (MathWorks, Inc.), and the experiment was presented using the online platform Testable (https://www.testable.org). The stimuli were Gabor patterns with different orientations, sinusoidal gratings (wavelength: 1.5°) at maximum contrast convolved with a two-dimensional Gaussian (standard deviation: 1°). We used three main conditions with orientations centred on 45°, 90°, or 135°. The orientation of test patterns could differ from these main orientations by 0°, ±5°, ±10°, ±15°, or ±20°. All stimuli were presented on a uniform grey screen.

Procedure

Participants sat in front of a computer in a small testing cubicle, at a distance of 75 cm from the screen (model: Dell OptiPlex 5260 aio series, size: 26.8 cm by 47.6 cm, 21.5 inch, resolution: 1920 x 1080 pixels). There were 27 unique trial types: nine orientation offsets for each of the three main orientations. Each trial type was presented 20 times in the experiment in interleaved order. Therefore, there were 540 trials in the actual experiment. In addition, 12 trials at the beginning of the task were practice trials that were not used for analysis.

Each trial contained two stimuli. The first grating was the reference, presented for 500 ms. This could be one of the three main orientations. This was followed by a 300 ms inter-stimulus interval during which only a grey screen and the fixation cross were presented. The second grating appeared on the screen until the participant responded. It could be tilted anticlockwise, clockwise, or remain unchanged relative to the first grating (Figure S1). Participants indicated via the keyboard left or right arrow keys whether the second grating was tilted anticlockwise or clockwise, respectively. They were instructed to respond as fast and as accurately as possible. If they were unsure about the change between the two stimuli, they must nevertheless press one key to continue the task. Throughout the run, participants were asked to fixate on a small, black cross in the middle of the screen. After the participant's response, a blank screen was presented for a 500 ms inter-trial interval.

During the 12 initial practice trials, participants could only choose the correct answer to move on to the next trial. The stimuli in the practice were easy to discriminate, including only orientation offsets of ±10° and ±20° for each of the three main orientations. Most participants could select the correct answer easily.

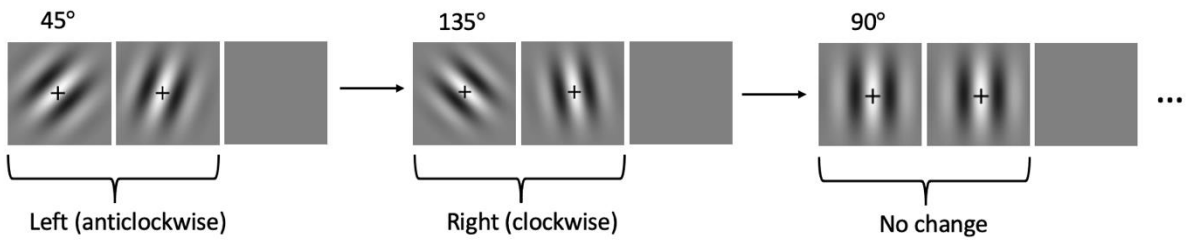

Figure S1. Experimental design. *Trial sequence.* The second (test) grating was either tilted anticlockwise or clockwise or was unchanged relative to the first (reference) grating. There were three main orientations around which the test grating could vary. The oblique effect is characterised by superior performance around cardinal (in this case, vertical) orientations. Note that a 500 ms inter-trial interval followed the observer's response. There was also a 300 ms inter-stimulus interval between the two gratings in each trial, not shown here. The second grating stayed on the screen until the observer responded.

Data analysis

We then determined the participant's orientation discrimination performance for each of the three main orientations. For each orientation offset, we calculated the proportion of trials that they chose the test stimulus as appearing more anticlockwise. Then we fit a cumulative Gaussian psychometric function to these data using a linear mean squares approach. The bandwidth (σ) of this Gaussian determines the slope of the curve. Narrower bandwidth corresponds to steeper slopes, meaning better performance (Figure S2A). We calculated two contrasts of interest. For determining the oblique effect (H_1) we subtracted the bandwidth for 90° main orientation from that for 45°. To estimate an assumed null effect (H_0), we subtracted the bandwidth for 135° from that for 45°.

We also determined the significance of these contrasts for each observer to use these results in our small-N approach. To this end, we bootstrapped the psychometric curve fits. We resampled the choices for each orientation offset with replacement, recalculated the choice proportions, and refit the psychometric curves 10,000 times. We then determined a one-tailed significance value by calculating the proportion of resamples for which the bootstrapped distribution of the two contrasts were ≤0.

Results

Psychometric curves for three representative observers are shown in Figure S2A. It is immediately apparent that the curve for 90° main orientation (red functions) has a steeper slope in all observers, that is a narrower bandwidth. We plot the contrasts between bandwidths for H_1 (45°-90°) and H_0 (45°-135°) in Figure S2B. All observers showed positive contrasts for H_1 as expected, while contrasts for H_0 ranged between positive and negative but most were in fact negative. Only one observer failed to reach significance on individual fits for H_1 and only one was significant for H_0.

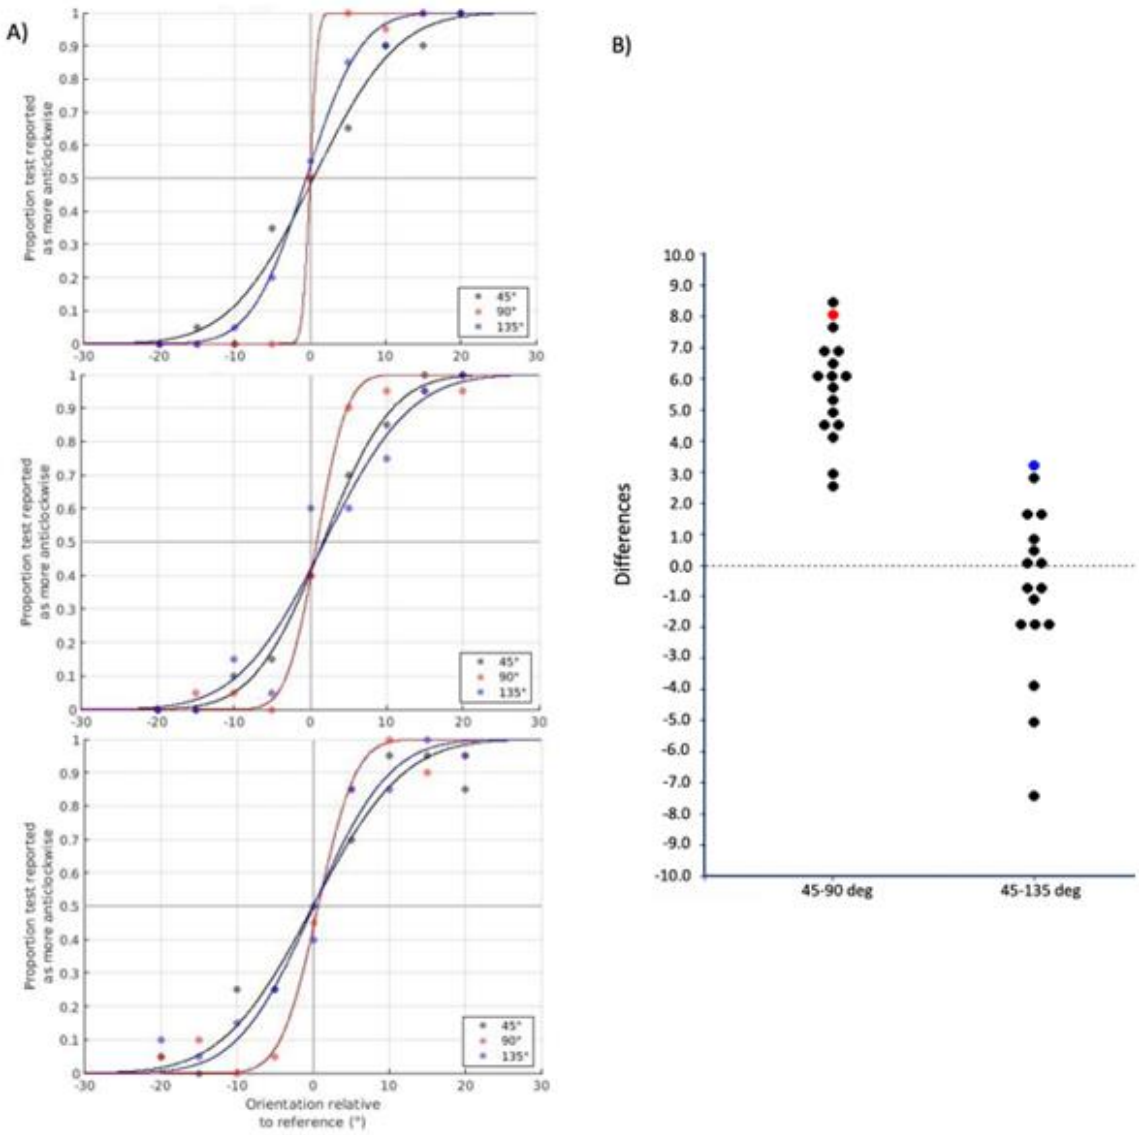

Figure S2. Oblique effect results. **A)**. Psychometric curves from three representative observers. The proportion of trials the test grating was seen as more anticlockwise is plotted against the orientation offset, separately for the three main orientations. Black: 45°, Red: 90°, Blue: 135°. **B)**. Contrasts across participants. Each dot denotes the difference in bandwidth between the conditions shown. The red dot indicates the observer who failed to reach significance for the contrast 45-90°. Conversely, the blue dot indicates the observer who showed a significant difference for the contrast 45-135°.