THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Power calculations in single-case neuropsychology

Citation for published version:

McIntosh, R & Rittmo, J 2021, 'Power calculations in single-case neuropsychology: A practical primer', *Cortex*, vol. 135, pp. 146-158. <https://doi.org/10.1016/j.cortex.2020.11.005>

Digital Object Identifier (DOI):

[10.1016/j.cortex.2020.11.005](https://doi.org/10.1016/j.cortex.2020.11.005)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Cortex

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Power calculations in single-case neuropsychology: a practical primer

Robert D McIntosh^{a,*}, Jonathan Ö Rittmo^a

^a*Human Cognitive Neuroscience, Psychology, University of Edinburgh, UK*

Abstract

Researchers and clinicians in neuropsychology often compare individual patients against healthy control samples, to quantify evidence for cognitive-behavioural deficits and dissociations. Statistical methods for these comparisons have been developed that control Type I (false positive) errors effectively. However, remarkably little attention has been given to the power of these tests. In this practical primer, we describe, in minimally technical terms, the origins and limits of power for case-control comparisons. We argue that power calculations can play useful roles in single-case study design and interpretation, and we make suggestions for optimising power in practice. As well as providing figures, tables and tools for estimating the power of case-control comparisons, we hope to assist researchers in setting realistic expectations for what such tests can achieve in general.

Keywords: Single-case; deficit; dissociation; power; statistical methods.

1. Introduction

In the past 25 years, a range of related statistical methods have been developed to compare the test scores of single-cases to those of matched control samples (e.g. Crawford & Howell, 1998; Crawford et al., 2011; Crawford & Garthwaite, 2005, 2007; McIntosh & Brooks, 2011). These case-control comparisons have put the classical concepts of neuropsychological deficit and dissociation onto a firmer quantitative footing (Crawford et al., 2003; McIntosh, 2018). They have acquired the status of gold standard inferential tests for investigating patterns

*Corresponding Author

Email addresses: r.d.mcintosh@ed.ac.uk (Robert D McIntosh), j.rittm@gmail.com (Jonathan Ö Rittmo)

of neuropsychological impairments at the single-case level, and they have broad applicability for defining diagnostic criteria in clinical settings.

The main focus for the development of these tests has been to provide transparent control over the rate of false positives (Type I or α errors), the frequency with which the test will return a spurious significant result when no true deficit (or dissociation) is present. But we should also be concerned about the rate of false negatives (Type II or β errors), the frequency with which the test fails to return a significant result when a true deficit (or dissociation) is present. To the extent that we can reduce the false negative rate, our test will have higher power (or sensitivity) to detect true effects. Even if we are unable to fully control the false negative rate, it is still relevant to estimate the likely power of our test, to inform study design and interpretation.

Remarkably little attention has been given to the issue of power in single-case studies. Of the three main papers that discuss power, one is concerned primarily with the effects of departures from normality (Crawford & Garthwaite, 2006b), the others with the consequences of choosing different operational criteria to classify dissociations (Crawford & Garthwaite, 2006a; McIntosh, 2018). Despite a growing awareness of the wider importance of statistical power in study design and interpretation, there is yet no general primer for power in case-control comparisons. The purpose of the present paper is to provide such a primer.

We aim to describe, in minimally technical terms, the nature and limits of power in testing for deficits and dissociations. Some of the points we make may seem obvious, even self-evident on reflection, especially to statistically-sophisticated readers. However, it is our impression that many researchers and practitioners who use case-control comparisons have never been prompted to consider these issues deeply, and may be unsure how (or why) to calculate power for these tests.¹ The figures and tables in this paper will be broadly sufficient to estimate power when testing for deficits and dissociations. We also refer the reader to the `singcar` package for the R statistical software environment (Rittmo & McIntosh, 2020), which implements a wider range of single-case tests (including Bayesian extensions of the frequentist tests covered in this primer), and associated power calculations (see [Appendix I](#) for details).

¹The first author admits to using these statistical tests for many years without having thought seriously about their probable power, or its determining factors.

2. Case-control comparisons as a form of outlier detection

In testing for a neuropsychological abnormality, we wish to estimate the probability that a patient’s performance for a given task or domain has been impacted by brain damage or dysfunction. Ideally, we would do this by evaluating current performance with respect to that person’s premorbid performance. This may indeed be possible if we wish to study the sequelae of a scheduled event such as neurosurgery, and we have had the foresight to characterise premorbid performance. But commonly the misfortune that has brought the patient to our attention is a sporadic event such as a stroke or head injury. If so, the best that we can usually do is to measure post-morbid performance, and compare this to the distribution of performance in a sample of healthy controls matched to the case on relevant characteristics (e.g. age, sex, handedness, years of education). This comparison aims to estimate the probability that we would observe performance as extreme as that of our patient if we were to sample a member of the healthy population at random. If we find that this probability is sufficiently low, we may infer that our patient does not belong to that population: we will have evidence for a neuropsychological abnormality.

Case-control comparisons are thus a form of outlier detection, but rather than targeting the most extreme observation in a sample (e.g. Grubbs, 1969), we target an observation (case) of interest and ask whether it is an outlier with respect to a separate (control) sample. In the standard normal distribution (the Z-distribution), only 5% of cases fall more than 1.96 standard deviations from the mean (2.5% in either tail) (Figure 1a). To apply this to the evaluation of a single-case, we could express the patient’s test score as a Z-score [i.e. (patient score - control mean)/control standard deviation)], and regard a Z-score lower than -1.96 or higher than 1.96 as evidence of abnormality. This would constitute a two-tailed test of abnormality with a significance criterion of .05. However, when studying neuropsychological cases, we are often interested specifically in performance that is lower than normal, indicating a deficit. To frame a one-tailed test, we could regard a Z-score below -1.65 as evidence of a deficit, because fewer than 5% of cases fall more than 1.65 standard deviations below the mean (Figure 1a).

Figure 1: **(a)** The standard normal (Z) distribution, showing two-tailed (upper and lower) cut-offs, and the one-tailed lower cut-off for a case-control test of deficit, with a significance (α) criterion of .05. **(b)** Two hypothetical patients (A and B), both of whom sustain a 2 SD performance deficit due to brain damage. The letters show the levels of pre-morbid performance, and the leftward arrows show the shift in performance due to the deficit. Because of differing levels of premorbid performance, patient A is cast into the abnormal range but patient B is not. Pre-morbid performance would need to be within the shaded region of the curve in order for a deficit of this size to be detected by this test. The power of the test is thus equal to the probability that the premorbid performance was in the shaded region of the Z-distribution. **(c)** Power of this Z-based test to detect different sizes of deficit, given a one-tailed significance (α) criterion of .05 (black line), or more conservative (.025) or liberal (.10) significance criteria (grey lines). Because this figure is derived from the Z-distribution, it shows the power that could be achieved with an unrestricted control sample. It therefore traces the upper limits on power for case-control tests of deficit.

These Z-based methods would be appropriate if we were to compare the patient against a large enough control sample (e.g. $n \geq 50$), because the sampling distribution of the statistic of interest would be approximately normal.² With more restricted control samples, however, using the Z-distribution would underestimate the likelihood of extreme scores, and inflate the rate of false positives. To avoid this problem, case-control tests based on the t-distribution have been developed, which take sample size into account. We will turn to these methods in Section 4, but for initial purposes the Z-score method provides a maximally simple framework within

²The minimum sample size for which we could use Z-scores is a somewhat arbitrary and context-dependent judgement. Here, we follow Crawford and Howell (1998) in suggesting that the use of Z-scores would be unproblematic given a control sample size of $n \geq 50$.

which to illustrate the origins and limits of power in case-control comparisons.

3. The origins and limits of power in case-control comparisons

The power of any hypothesis-testing procedure is the probability that it will correctly reject the null hypothesis when the null hypothesis is false or, more directly, that it will detect a true effect when present. For a case-control test of deficit, for instance, the power of the test would be its ability to detect a true deficit when present. Power is the simple complement of the false negative rate (β), and is given by $1-\beta$, though sometimes expressed as a percentage $[(1-\beta)*100]$.

If a patient who was pre-morbidly a part of the healthy population incurs a deficit in a tested ability, their score will be shifted relative to controls. The size of this shift is the effect of the deficit, which we can also express in terms of standard deviations of the control mean. This gives us a standardised effect size of the deficit, which Crawford et al. (2010) labelled Z_{CC} (i.e. the “case-controls” Z-score). We cannot measure the deficit directly, however, we can only observe the post-morbid performance, which is co-determined by the level of premorbid performance and the size of the deficit. Whether or not a given cut-off will detect a deficit of a given size when evaluating post-morbid performance thus depends critically upon the patient’s (unobserved) premorbid ability.

Two examples are depicted in Figure 1b, with the cut-off for a deficit set at 1.65 SD below the mean. These data are invented, so we can have full knowledge of both the size of the deficit and the pre-morbid ability. Patient A performed pre-morbidly at the mean level of controls, but acquired a 2 SD deficit due to brain damage so would be shifted to an observed post-morbid Z-score of -2 ; this is below the cut-off, so the deficit would be detected. But Patient B, who pre-morbidly scored 1 SD above the mean, and sustained exactly the same deficit, would be shifted to a Z-score of -1 , and the deficit would not be detected. In general, for a deficit of 2 SD to be detected, a patient’s premorbid performance would need to have been within the shaded region in Figure 1b. Premorbid performances higher than this would not be pushed into the abnormal range by a 2 SD deficit.

The power of a one-tailed test to detect a 2 SD deficit is therefore the same as the probability that the patient’s premorbid performance was within the shaded portion of the distribution up to .35 standard deviations above the mean (2 SD above the cut-off of -1.65). This is given by the area under the standard normal curve below .35, which can be obtained from a relevant Z-table or calculator: 63.7%

of observations fall less than .35 SDs above the mean, so our test has .64 power to detect a deficit of effect size 2 SD. If instead the deficit were 1.65 SD, the shaded zone would perfectly cover the lower half of the distribution, and the power of our test would be exactly .50. Power for a case-control test of deficit is conceptually simple: *it is the probability that the (unobserved) pre-morbid performance of the patient was no more than the deficit size above the cut-off for a deficit.*

Figure 1c maps the relationship between deficit size and power for a one-tailed test of deficit based on the Z-distribution, using a conventional significance (alpha) criterion of .05 (equivalent curves for alphas of .10 and .025 are also shown for comparison). There are two important points to note. The first, which is obvious from the figure, is that high power ($> .80$) is achieved only for very large effect sizes, greater than 2.5 SD. The second is that, because this figure is derived from the Z-distribution, it shows the power that could be achieved with an unrestricted control sample. It therefore traces the upper limits on power for case-control tests of deficit.

4. Power in testing for deficits

The Z-score method would be appropriate if the control sample were large enough (e.g. $n \geq 50$) that the sampling distribution closely approximates the normal distribution. But control samples in single-case studies are usually much smaller than this (sometimes even smaller than 10), and for such restricted samples one must use the family of t-distributions. These resemble the Z-distribution, but are somewhat flatter, with relatively more observations in the tails, so that the likelihood of observing an extreme score (say, 2 SD below the mean) is relatively higher. A test of deficit based upon the Z-distribution would overestimate the rarity of extreme scores in a restricted sample, increasing the rate of false positives. Crawford & Howell (1998) thus proposed a modified t-test for the detection of single-case deficits. This method controls Type I error rate appropriately with restricted control samples, and its results converge with those of the Z-score method when control samples are large. Crawford and Howell's modified t-test is now the gold standard Test of Deficit (TD) in neuropsychology.

Figure 2: Power of the case-control Test of Deficit (TD: Crawford & Howell, 1998) to detect different effect sizes of deficit (Z_{CC}) in a single case, with control sample sizes ranging from 4-50, assuming a one-tailed significance criterion of .05. Tabulated values for one and two-tailed tests of deficit are shown in Table A1. Note that the Bayesian Test of Deficit (Crawford et al., 2010) produces equivalent results to the TD, and has frequentist properties, so may be considered to have the same effective levels of power. Note also that the Unstandardised Difference Test (UDT) (Crawford & Garthwaite, 2005) is identical to applying the TD to between-task discrepancy scores. Power for the UDT is thus identical with power the TD, but related to the effect size of a between-task discrepancy (Z_{DCC}) which depends on the intertask correlation, rather than the effect size of a deficit (Z_{CC}).

The t-distribution is more complex than the standard normal which is reflected in power calculations for the TD (see Appendix II for details). Not only is power restricted by the pre-morbid ability of the patient but also by the control sample size. However, power for the TD is conceptually similar to that of the Z-score method in that *the upper limit of power is the probability that the (unobserved) pre-morbid performance of the patient was no more than the deficit size above the cut-off for a deficit*. Figure 2 shows how sample size affects the power of the (one-tailed) TD, for a range of (large) effect sizes (see Table A1 for tabulated values, including two-tailed tests). Power rises steeply as the control sample size increases within the very small range (< 10), but asymptotes rapidly, with severely diminishing returns above a sample size of around 16-24. Beyond this range, power tends towards the

maximum possible for this test, already mapped in Figure 1c. This ceiling level of power is inherently low, so the test of deficit is only really applicable to the study of effect sizes that would be very large in almost any other context ($Z_{CC} \geq 2$). High power ($\geq .80$) is inherently unachievable for effect sizes of 2.5 or less, and very high power ($\geq .90$) can only be expected for extreme deficits ($Z_{CC} > 3$), equivalent to an individual’s performance being decremented by around three-quarters of the full range of abilities in the healthy population. Neuropsychological case-studies are only viable at all because brain lesions can have such large effects upon cognitive-behavioural abilities.

5. Testing for dissociations

Neuropsychologists are often less interested in whether a patient has a deficit for a task (X), than in whether performance of that task (X) is abnormally impaired with respect to some other task (Y). Statistically, we wish to evaluate whether the between-task discrepancy is more extreme in the single case than in the healthy population. Crawford and colleagues have provided two methods based on the t-distribution (or asymptotic approximations of it) to make this comparison, whilst controlling the false positive rate appropriately (Crawford & Garthwaite, 2005). The simplest solution is applicable only when tasks X and Y are both scored on the same scale, with equivalent between-subject variances in the healthy population. In this special case, the test can be done directly on the paired XY differences, which is the basis of the Unstandardised Difference Test (UDT).³ More often, however, the two tasks will be on different scales so the scores must be standardised (expressed as Z-scores) before the differences are calculated, which is the basis of the Revised Standardised Difference Test (RSDT). The RSDT is marginally less powerful than the UDT but much more widely applicable, so we will focus exclusively on the RSDT here.

The patient may present with an abnormal discrepancy between tasks if brain damage has affected ability in one domain (tapped by task X) but not another

³As a rule-of-thumb, the UDT may be applicable to pairs of tasks for which it would be sensible to perform a paired t-test within the control group; for instance, to compare performance on the same reaching task performed under monocular and binocular viewing conditions. Note that the UDT is equivalent to applying the TD to the between-task discrepancy scores, rather than to the scores of a single task. Power for the UDT is thus identical with power for the TD (Figure 2 and Table A1), but related to the effect size of a between-task discrepancy (Z_{DCC} ; see Section 6) rather than the effect size of a deficit (Z_{CC}).

(tapped by task Y), or has affected one domain more severely than the other. These two scenarios, of selective and differential deficit, map onto a theoretical distinction between a classical dissociation and a strong dissociation (Crawford et al., 2003; Shallice, 1988). A classical dissociation is held to have special status, licensing stronger inferences about the fractionation of underlying mental structure. Statistically, however, a classical dissociation requires the assertion of the null hypothesis of no deficit for the less affected task. But the severe limits on single-case power (Figure 1c) make it clear that we can never confidently rule out the possibility that a deficit as large as 2.5 SD has been missed, let alone exclude non-trivial deficits smaller than this. Consequently, we can never confidently distinguish a classical from a strong dissociation (Crawford et al., 2003; McIntosh, 2018). In light of this limitation, Crawford et al. (2003) have suggested that classical dissociations should be renamed, “putatively classical”. More recently, McIntosh (2018) has argued that the distinction should be abandoned for practical purposes, and that a significant intertask discrepancy should be the sole criterion for establishing a dissociation. Following this logic, we will consider the RSDT (and its relatives) as constituting a simple test of dissociation.

6. The importance of intertask correlation in testing for dissociations

A dissociation might involve a deficit of 2 SD on task X with no deficit on task Y, and an equivalently-sized dissociation would arise if a deficit of 3 SD on task X were accompanied by a milder deficit of 1 SD on task Y. For ease of exposition, let us imagine that these two patterns of dissociation have been sustained by two patients who both happened to have performed pre-morbidly at exactly the control mean level on both tasks. This situation is depicted in Figure 3, for the two patients and a large matched control group ($n = 50$). The dissociation between tasks is not captured by the score on task X or Y alone, but by the discrepancy between them ($Y-X$), represented graphically by the distance from the clockwise diagonal.

Figure 3: **(a)** Standardised performance on two tasks, X and Y, for 50 healthy control participants (open circles) and two hypothetical patients (black and grey circles). The size of the intertask discrepancy is represented by the distance from the clockwise diagonal. Both patients have an equivalent discrepancy score ($Y-X$) of 2, but the abnormality of this discrepancy depends on how extreme it is with respect to the distribution of discrepancies in the control group. The three panels illustrate that when the two tasks are more highly correlated in the control sample, the spread of discrepancies in controls is reduced, so the discrepancies in the patients are relatively more extreme. That is, the standardised effect size of the discrepancy (Z_{DCC}) is increased, as shown in panel b. **(b)** The relationship between discrepancy score and discrepancy effect size (Z_{DCC}), mediated by the strength of the intertask correlation (r) in the control sample. This plot follows the function $Z_{DCC} = (Z_{CC_Y} - Z_{CC_X}) / \sqrt{2 - 2r}$, where Z_{CC_X} is the standardised deficit on task X and Z_{CC_Y} is the standardised deficit on task Y.

Our hypothetical patients both have a discrepancy of 2, so they are equidistant from this diagonal. The abnormality of this discrepancy depends on how extreme it is with respect to the distribution of discrepancies in the control group (the spread of unfilled circles' distances from this diagonal). The three panels of Figure 3a illustrate that when the two tasks are more highly correlated, the spread of discrepancies in controls is reduced, so the discrepancies in our patients are relatively more extreme. This can be expressed by dividing the discrepancy score by the standard deviation of the discrepancies in the control group, yielding a standardised effect size for the discrepancy between tasks (Z_{DCC} : Crawford et al., 2010). When this effect size is larger, the discrepancy is more abnormal with respect to the control sample: there is stronger evidence for a dissociation. Intuitively, the more strongly-related two tasks are in the population, the more surprised we should be by a substantial discrepancy in a single case.

This influence of intertask correlation is easy to capture mathematically. The standardised control scores for each task have (by definition) a mean of zero and a standard deviation of 1, so the mean of the distribution of discrepancies is always zero, but the standard deviation of the discrepancies is given by $\sqrt{2 - 2r}$ where r is the correlation between tasks. Only when the intertask correlation is .50 will the standard deviation of discrepancies be equal to 1. In this case, the standardised effect size of the discrepancy (Z_{DCC}), obtained by dividing the discrepancy score by the standard deviation of discrepancies, will equal the discrepancy score itself (2 in our example). However, if the intertask correlation is less than .50, the standard deviation of discrepancies will be greater than 1 so Z_{DCC} will be smaller; and if the correlation exceeds .50, the standard deviation of discrepancies will be less than 1 so Z_{DCC} will be larger. The relationship between the intertask discrepancy score and Z_{DCC} , governed by the strength of the intertask correlation, is plotted in Figure 3b for four different sizes of discrepancy. This relationship implies that pairs of tasks that are more highly correlated in healthy populations provide the potential for more highly powered tests of dissociation.⁴

⁴Albeit it is possible that abilities that are highly correlated in the healthy brain might also be less likely to give rise to substantial dissociations after brain damage.

7. Power in testing for dissociations

A case-control test of dissociation (RSDT and its relatives) is a matter of estimating the proportion of the healthy population who would produce an intertask discrepancy more extreme than that of the single-case. The conceptual basis of power for this test is similar to that for the TD: *the upper limit of power is the probability that the patient's (unobserved) pre-morbid discrepancy between tasks was no more than the acquired discrepancy size away from the cut-off for a dissociation*. It is less tractable to estimate power analytically for the RSDT than for the TD, because its test statistic is only an approximation to the t-distribution (Crawford & Garthwaite, 2005; Garthwaite & Crawford, 2004), so a simulation approach has been used to generate the power curves shown in Figure 4 (see Appendix II for details). Note that the power functions in Figure 4 are for a two-tailed RSDT, which is the default choice if the researcher has no prior commitment to the direction of a dissociation in a given patient (e.g. Crawford & Garthwaite, 2005; Crawford et al., 2003). However, in practice, we often do have a strong prediction based on a patient's prior behaviour and/or pattern of brain damage (McIntosh, 2018). If so, then one-tailed tests of dissociation should be considered, for the higher power they offer (see Section 10).

Figure 4: **(a)** Power of the Revised Standardised Difference Test (RSDT: Crawford & Garthwaite, 2005) to detect different intertask discrepancies, mediated by the strength of the intertask correlation (r) in the control sample, assuming a two-tailed significance criterion of .05, and a very large control sample ($n = 100$). **(b)** Power of the RSDT to detect different effect sizes of discrepancy (Z_{DCC}) in a single case, with control sample sizes ranging from 4-50, assuming a one-tailed significance criterion of .05. Tabulated values for the one and two-tailed RSDT are shown in Table A2.

Figure 4a confirms that the power of the (two-tailed) RSDT to detect different levels of intertask discrepancy in a single case is modulated dramatically by the intertask correlation in healthy participants. This is just an expression of the dynamic already shown in Figure 3b: for any given discrepancy, the RSDT gains power with intertask correlation (Figure 4a) because the standardised effect size of the discrepancy (Z_{DCC}) increases (Figure 3b). A more general way to investigate the power of the RSDT is thus in terms of its relationship with Z_{DCC} , which takes intertask correlation into account. Figure 4b shows how sample size affects the power of the (two-tailed) RSDT, for a range of (large) discrepancy effect sizes (see Table A2 for tabulated values, including one-tailed tests). These functions are qualitatively similar to those for a test of deficit (Figure 2), with high power achievable only for very large effect sizes, and control sample size providing little additional benefit after the first 16-24 participants.

8. Why sample size doesn't help (beyond a point)

Figures 2 and 4b show that sample size is a powerful determinant of power when the control group is small (<16), but they also show that it is not the fundamental limiting factor. There are hard limits on power, which do not apply to more typical hypothesis tests, such as paired or independent t-tests. For those more typical tests, the critical comparison can always be boiled down to testing whether the mean difference between two conditions or groups is different from zero. The precision with which this judgement can be made depends upon how close the estimate of the mean difference is likely to be to the true (population) mean difference. This is indexed by the standard error of the mean, which is inversely proportional to the square root of the sample size ($SEM = SD/\sqrt{N}$), so that increasing the sample size reduces the standard error, allowing us to discriminate smaller and smaller differences from zero. In principle, whatever size of effect we are targeting, any level of power (up to 100%) can be achieved if only we have a large enough sample size.

Case-control comparisons, by contrast, do not boil down to a comparison between two means, but to a comparison of one observation to the spread of observations in a sample: they are a targeted form of outlier detection (Section 2; Figure 1). The most relevant measure of spread here is not the standard error, which characterises the sampling distribution of the mean, but the standard deviation, which characterises the distribution of observations in the population. Unlike the standard error, the standard deviation does not get smaller as we increase the sample size, it just gives

a more accurate estimate of the true population value. If the control sample is large enough (e.g. $n \geq 50$), then its mean and standard deviation will be very good estimates of the population values, and the power of our case-control test will be maximal, but not necessarily high (Figure 1c). *Power in case-control comparisons is always limited by the probability that the (unobserved) pre-morbid performance of the patient was no more than the deficit/discrepancy size away from the critical cut-off.* For instance, if our patient was pre-morbidly at the top end of the control range of abilities, we will inevitably fail to detect all but the most severe of acquired deficits, regardless of how many control participants we test.

9. Using power calculations in single-case studies

The canonical role of a priori power calculations is to help a researcher set an appropriate sample size to achieve a desired (high) level of power for a hypothesis test. But the foregoing discussion shows that the power of case-control comparisons is chronically limited. Beyond testing a reasonably-sized control sample (e.g. $n > 16$), there may be little that the researcher can do to increase power (though see Section 10). This suggests that we will rarely be able to use a priori power calculations to ensure that a case-control study is adequately powered by design. A more common role for these calculations may be in helping researchers to estimate how sensitive their single-case tests will be to the effects of interest. This might be done with respect to prior estimates of expected effect sizes, or it might be done without such knowledge. Brief examples of the use of power calculations in both scenarios will be given, from recent neuropsychological studies.

Hassan et al. (2020) re-examined the claim that patient DF, a much-studied patient with visual form agnosia following bilateral lesions to the ventral visual stream, does not experience a normal size-weight illusion. The size-weight illusion is a striking phenomenon whereby smaller objects are perceived as heavier during lifting than larger objects of equal weight (Buckingham, 2014). This illusion implies that perceptions of heaviness are influenced by sensory cues about object size. The possible disruption of this illusion in patient DF suggests that the critical size cues may be processed in the ventral visual stream. The specific claim was that DF does not experience the size-weight illusion when object size is cued visually, but has a normal strength illusion when kinaesthetic size cues are provided (Dijkerman et al., 2004; McIntosh, 2000). However, the evidence in these prior reports was informal, and involved no comparisons with healthy controls, so the claim was in need of

confirmatory testing.

A first step was taken in a pilot study with patient DF, to estimate her visual size-weight illusion. This pilot study suggested that DF had a very large deficit in illusion magnitude, with an effect size (Z_{CC}) of -3.4 relative to a small group of controls ($n = 6$). A formal test of this deficit was then mounted, with a much larger age-matched control sample ($n = 28$), and a parallel assessment of the kinaesthetic size-weight illusion. Even using a more modest estimate for the size of the deficit ($Z_{CC_x} = -3$), the a priori power to detect it using a one-tailed alpha of .05 approached .90 (see Figure 2 or Table A1). Since the theory predicted that the kinaesthetic size-weight illusion should be unaffected ($Z_{CC_y} = 0$), we were also able to define an expected discrepancy of 3 between the visual and kinaesthetic conditions. Moreover, the correlation between these conditions in the control group was $r = .62$, so that the estimated effect size for the dissociation was even larger than the discrepancy score ($Z_{DCC} = 3.44$) (see Figure 4a), and the a priori power for a directional (one-tailed) test of dissociation (RSDT) exceeded .90 (see Table A2). This is a (possibly rather rare) instance in which the tests of deficit and dissociation were estimated to be highly powered. It is also, to our knowledge, the first neuropsychological single-case study to feature a priori power calculations.⁵

Another, maybe more routine use for power calculations is in setting realistic expectations for the range of abnormalities to which our case-control tests will be sensitive. Mitchell et al. (2020) published a protocol for the assessment of visually-guided reaching deficits in patients with mild-to-moderate Alzheimer’s disease and patients with Mild Cognitive Impairment. A control sample size of 24 was proposed in order to ensure that individual tests of deficit would achieve close to the maximum possible power (see Figure 2). Nonetheless, it was emphasised that the tests could only be expected to be sensitive to large abnormalities, with less than .80 power for deficit effect sizes below around 2.5 (see Figure 2 or Table A1). As in this multiple single-case study, power calculations can inform the choice of research question, the design of the study, and the interpretation of outcomes, even if the practical limits on power cannot be overcome.

⁵The experiment confirmed the expected deficit in the visual size weight illusion, albeit with a more modest effect size estimate: $Z_{CC} = -1.76$ [95% $CI : -2.345, -1.155$]; but the dissociation between visual and kinaesthetic conditions was not confirmed by the RSDT (Hassan et al., 2020). Because the two conditions were parallel versions of the same task, it might have been legitimate to use the UDT rather than the RSDT to test this dissociation (but this would not have found a dissociation either).

10. Optimising power in single-case studies

Low power in single-case studies is aggravated by small control samples, but cannot be wholly ameliorated by increasing control numbers (Section 8). Beyond sample size, the primary determinants of statistical power are the alpha significance criterion, and the size of the effect targeted. This suggests a few possible strategies for optimising the power of these tests.

The simplest way to increase the power of any significance test is just to adopt a less stringent significance criterion, for instance $p < .10$ instead of the more conventional $p < .05$ (see Figure 1c). This would raise the power of the test to identify a true effect when present, though at the cost of increasing the rate of false positives when no true effect is present (from 5% to 10%). Whether this is a worthwhile trade-off depends on how costly the researcher considers false negative and false positive errors to be in the context of their study. A widely-known suggestion, from Cohen (1988), is that a study using a significance criterion of .05 (5% false positive rate) might aim for a power of .80 (20% false negative rate), implying that false positives are four times as costly as false negatives. This might be a reasonable suggestion, assuming that the researcher has selected the conventional (but arbitrary) significance criterion of .05 proposed by Fisher (1925). However, neither Fisher nor Cohen were laying down rules to be followed blindly; they were encouraging researchers to choose criteria appropriate to the aims of their own particular study.

Given that the power for a case-control comparison to detect an effect size of 2 barely exceeds .60 (~40% false negative rate), the only way to preserve a four-to-one balance of false negatives to false positives, if desired, would be to relax the significance criterion to .10. It might indeed be rational to relax the significance criterion (perhaps even further than .10) for some purposes, such as when screening for relatively common conditions, or in exploratory research, in order to increase the power to detect possible abnormalities. On the other hand, or in other contexts, researchers (or reviewers) may be unwilling to tolerate the elevated risk of false positives that this would entail. Another way to maximise power is to adopt a one-tailed significance criterion if this is reasonable in the context of the study; specifically, when the researcher has a directional prediction, and no theoretical interest in interpreting an effect in the opposite direction. One-tailed tests are the default choice (by definition) for a test of deficit, but we may also have a prior commitment to a directional dissociation, in which case a one-tailed test

of dissociation also may be appropriate. An example of this was provided in the size-weight illusion study discussed in Section 9 (Hassan et al., 2020).

Other strategies might focus on the targeted effect size. This primer has shown that case-control comparisons can achieve respectable power ($> .60$) only for effects that would be very large in most other contexts (standardised effect sizes of 2 and upward). A first consideration, then, is that researchers should target large effects, and accept that more subtle effects cannot be studied reliably. As often noted, neuropsychological case-studies are viable precisely because brain damage and dysfunction can have such pronounced effects upon cognition and behaviour. A second consideration is that, although we cannot alter the neuropsychological consequences of a patient's lesion, we can try to limit the background noise against which it is measured. Reducing measurement error will give a more precise estimate of the patient's post-morbid performance, enhancing our ability to identify true effects of interest. Similar improvements in the precision of estimation within control participants will reduce the control sample variance, thereby increasing the standardised effect size associated with a given deficit, improving our power to detect it.

Strategies for reducing measurement noise are not different in single-case studies than in other behavioural studies. We should aim to select and develop valid, reliable tests of a targeted ability, with scores that are well-distributed in healthy samples, and not prone to floor or ceiling effects (Crawford & Garthwaite, 2006b). We should collect more observations per participant, if possible, and make them as systematically as we can, for more reliable estimates of performance. Case-control comparisons also allow for the inclusion of covariates, providing another way to reduce residual variance (Crawford et al., 2011). One degree of freedom is sacrificed for each covariate included, so we should build this into our design by testing an additional control participant for each covariate we intend to include. The burden of extra recruitment is offset somewhat by greater flexibility, since controls no longer need to be matched so closely to the patient, but can bracket the patient on the covariate of interest (Crawford et al., 2011).

The present focus is on the power of specific comparisons, but it should be noted that a convincing neuropsychological dissociation requires more than a one-off contrast between two tasks (Shallice, 1988). An optimal design for a detailed single-case study would include multiple tasks to triangulate the cognitive functions of interest, performed on more than one occasion. The integration of evidence across tasks and replications requires informed, qualitative evaluations, and might also be

aided by multivariate extensions of case-control comparisons, yet to be developed.

11. Conclusion

Our first purpose in this primer has been to present and explain the main facts of statistical power for case-control comparisons in neuropsychology. The most salient fact is that power is chronically low, as has been noted in prior discussions (e.g. Crawford & Garthwaite, 2006a, 2006b; McIntosh, 2018). The limiting factor is the nature of the statistical test itself, which we have characterised as a form of outlier detection, but the problem is aggravated by the use of small control samples. These tests have been designed to control false positive rates even with very small control samples, but a moment's consideration of the likely false negative rates should warn against using such small samples. We would suggest that control sample sizes for single-case studies should not be below eight, and that two (or even three) times that many control participants is generally desirable. Some other recommendations for optimising power were given in Section 10.

Our second purpose has been to provide tools for power calculations when testing for deficits and dissociations. Figures 2 and 4b and Tables A1 and A2 show power for tests of deficit and dissociation across a range of likely parameters, and should enable ballpark estimates, adequate for most purposes. The `singcar` package for R supports power calculations, and for a wider range of tests (see Appendix I for details). However, the ability to calculate power for a given effect size may engender a false sense of precision or control. In fact, the estimates of effect size in single-case studies are themselves subject to a high degree of uncertainty, with very wide confidence (or credible) intervals, and power calculations based on such estimates will also be imprecise. Moreover, in a typical neuropsychological case study, we have little control over the effect sizes we encounter, and limited scope to achieve a desired level of power. Nonetheless, a priori power calculations can still play useful roles in study design and interpretation, as illustrated in Section 9.

Finally, in highlighting the practical realities of power in case-control comparisons, we hope to assist researchers in setting realistic expectations for what such tests can achieve in general. This brief primer may thus inform a broader discussion of the epistemological status of the quantitative neuropsychological case study as a method for investigating the organisation of mind and brain.

Open data and tools

Code used for calculations and simulations is available at github⁶. Power calculations for a wider range of single-case tests are supported by the R package `singcar` (Rittmo & McIntosh, 2020); see [Appendix I](#) for details.

References

- Buckingham, G. (2014). Getting a grip on heaviness perception: A review of weight illusions and their probable causes. *Experimental Brain Research*, 232(6), 1623–1629. <https://doi.org/10.1007/s00221-014-3926-9>
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences*. L. Erlbaum Associates. <https://books.google.se/books?id=gA04ngAACAAJ>
- Crawford, J., & Garthwaite, P. (2005). Testing for Suspected Impairments and Dissociations in Single-Case Studies in Neuropsychology: Evaluation of Alternatives Using Monte Carlo Simulations and Revised Tests for Dissociations. *Neuropsychology*, 19(3), 318–331. <https://doi.org/10.1037/0894-4105.19.3.318>
- Crawford, J., & Garthwaite, P. (2006a). Detecting dissociations in single-case studies: Type I errors, statistical power and the classical versus strong distinction. *Neuropsychologia*, 44(12), 2249–2258. <https://doi.org/10.1016/j.neuropsychologia.2006.05.019>
- Crawford, J., & Garthwaite, P. (2007). Comparison of a single case to a control or normative sample in neuropsychology: Development of a Bayesian approach. *Cognitive Neuropsychology*, 24(4), 343–372. <https://doi.org/10.1080/02643290701290146>
- Crawford, J., & Garthwaite, P. (2006b). Methods of testing for a deficit in single-case studies: Evaluation of statistical power by Monte Carlo simulation. *Cognitive Neuropsychology*, 23(6), 877–904. <https://doi.org/10.1080/02643290500538372>
- Crawford, J., Garthwaite, P., & Gray, C. (2003). Wanted: Fully Operational Definitions of Dissociations in Single-Case Studies. *Cortex*, 39(2), 357–370. [https://doi.org/10.1016/S0010-9452\(08\)70117-5](https://doi.org/10.1016/S0010-9452(08)70117-5)

⁶https://github.com/jorittmo/neuropsychology_power_primer

- Crawford, J., Garthwaite, P., & Porter, S. (2010). Point and interval estimates of effect sizes for the case-controls design in neuropsychology: Rationale, methods, implementations, and proposed reporting standards. *Cognitive Neuropsychology*, 27(3), 245–260. <https://doi.org/10.1080/02643294.2010.513967>
- Crawford, J., Garthwaite, P., & Ryan, K. (2011). Comparing a single case to a control sample: Testing for neuropsychological deficits and dissociations in the presence of covariates. *Cortex*, 47(10), 1166–1178. <https://doi.org/10.1016/j.cortex.2011.02.017>
- Crawford, J., & Howell, D. (1998). Comparing an Individual’s Test Score Against Norms Derived from Small Samples. *The Clinical Neuropsychologist*, 12(4), 482–486. <https://doi.org/10.1076/clin.12.4.482.7241>
- Dijkerman, H., Lê, S., Démonet, J.-F., & Milner, A. (2004). Visuomotor performance in a patient with visual agnosia due to an early lesion. *Cognitive Brain Research*, 20(1), 12–25. <https://doi.org/10.1016/j.cogbrainres.2003.12.007>
- Fisher, R. A. (1925). *Statistical Methods for Research Workers*. Oliver; Boyd. <https://books.google.se/books?id=I0NBAAAAIAAJ>
- Garthwaite, P., & Crawford, J. (2004). The distribution of the difference between two t -variates. *Biometrika*, 91(4), 987–994.
- Grubbs, F. E. (1969). Procedures for detecting outlying observations in samples. *Technometrics*, 11(1), 1–21.
- Harrison, D. A., & Brady, A. R. (2004). Sample Size and Power Calculations using the Noncentral t-distribution. *The Stata Journal*, 4(2), 142–153. <https://doi.org/10.1177/1536867X0400400205>
- Hassan, E. K., Sedda, A., Buckingham, G., & McIntosh, R. D. (2020). The size-weight illusion in visual form agnostic patient DF. *Neurocase*, 1–8. <https://doi.org/10.1080/13554794.2020.1800748>
- McIntosh, R. D. (2018). Simple dissociations for a higher-powered neuropsychology. *Cortex*, 103, 256–265. <https://doi.org/10.1016/j.cortex.2018.03.015>
- McIntosh, R. D. (2000). Seeing size and weight. *Trends in Cognitive Sciences*, 4(12), 442–444. [https://doi.org/10.1016/S1364-6613\(00\)01556-4](https://doi.org/10.1016/S1364-6613(00)01556-4)

- McIntosh, R. D., & Brooks, J. L. (2011). Current tests and trends in single-case neuropsychology. *Cortex*, 47(10), 1151–1159. <https://doi.org/10.1016/j.cortex.2011.08.005>
- Mitchell, A. G., McIntosh, R. D., Rossit, S., Hornberger, M., & Pal, S. (2020). Assessment of visually guided reaching in prodromal Alzheimer’s disease: A cross-sectional study protocol. *BMJ Open*, 10(6), e035021. <https://doi.org/10.1136/bmjopen-2019-035021>
- Rittmo, J. O., & McIntosh, R. D. (2020). singcar: Comparing Single Cases to Small Samples. R package version 0.1.1. <https://CRAN.R-project.org/package=singcar>
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge University Press.

Appendix I: Power calculations using singcar

The `singcar` package for R (Rittmo & McIntosh, 2020), available at <https://cran.r-project.org/web/packages/singcar>, implements the following single-case tests: Test of Deficit (Crawford & Howell, 1998); Unstandardised Difference Test (Crawford & Garthwaite, 2005); Revised Standardised Difference Test (Crawford & Garthwaite, 2005); Bayesian Test of Deficit (Crawford & Garthwaite, 2007); Bayesian Standardised Difference Test (Crawford & Garthwaite, 2007); Bayesian Test of Deficit with Covariates (Crawford et al., 2011); Bayesian Standardised Difference Test with Covariates (Crawford et al., 2011).

The `singcar` package also includes functions to perform power calculations for each of the above tests. The package is fully documented, with examples. We provide a brief guide here on how to use the package to calculate power for the Test of Deficit (TD) and the Revised Standardised Difference Test (RSDT).

First, start by installing and loading the package.

```
R> install.packages("singcar")
R> library(singcar)
```

To calculate power for the TD, use the `TD_power()` function. You must supply the expected case score and sample size. The `mean` and `sd` of the control sample can also be specified; if these are not specified then the default values of `mean` 0 and `sd` 1 are used, so that the expected case score is interpreted as a difference from the mean in standard deviations. If `alpha` is not specified, then a conventional

value of .05 is assumed by default. For a one-tailed test, specify the alternative “less” or “greater”. For a two-tailed test, specify the alternative “two.sided”. If the alternative is not specified, then a default value of “less” is assumed for `TD_power()`, giving a one-tailed test of deficit (assuming that a low score represents poorer performance). It is good practice to specify the alternative, in order to be explicit about the direction of the hypothesis being tested.

In the example below, the arguments supplied are the expected case score, the mean and sd of controls, and the control `sample_size`. The alternative “less” specifies a one-tailed test. The level of power is returned by the function (0.58196).

```
R> TD_power(case = 70,
+ mean = 100,
+ sd = 15,
+ sample_size = 16,
+ alternative = "less")
```

OUTPUT >

```
[1] 0.5819579
```

`TD_power()` can also be used to calculate the required sample size instead of power, by setting the desired level of power, and leaving `sample_size` blank. In the example below, we ask for the sample size required to achieve .6 power, given the case and control scores used in the previous example. The function reports that a power of 0.57806 can be obtained with a control sample size (n) of 15, but that power will not increase by more than 0.5% for any additional participant added after that (i.e. power is already close to the asymptotic maximum after 15 participants).

```
R> TD_power(case = 70,
+ mean = 100,
+ sd = 15,
+ power = .6,
+ alternative = "less")
```

OUTPUT >

Power (0.57806) will not increase more than 0.5% for any additional participant over n = 15

	n	power
1	15	0.5780555

Note that the algorithm stops when power does not increase more than 0.5% for any additional participant, even if this is below the desired level of power. However, a higher degree of specificity can be obtained by changing the `spec` parameter from its default value of `.005`. For instance, if we set the value of `spec` in the previous example to `.001`, then the function reports that a power of 0.60126 can be obtained with a control sample size (`n`) of 24.

```
R> TD_power(case = 70,
+ mean = 100,
+ sd = 15,
+ power = .6,
+ alternative = "less",
+ spec = .001)
```

OUTPUT >

```
  n    power
1 24 0.6012626
```

To calculate power for the RSDT, use the `RSDT_power()` function. You must supply the expected case scores for the two tasks, and the control sample size. The `mean` and `sd` of the control samples for the two tasks can also be specified; if these are not specified then the default values of `mean` 0 and `sd` 1 are used, so that the expected case scores are interpreted as a differences from the mean in standard deviations. `RSDT_power()` additionally requires an estimate of the inter-task correlation in the control sample (`r_ab`); if this is not specified then a correlation of `.5` is assumed by default. If `alpha` is not specified, then a conventional value of `.05` is assumed by default. For a one-tailed test, specify the alternative “`less`” or “`greater`”. Take care when changing this parameter for `RSDT_power()`, because the direction of the expected effect depends on which task is set as A and which is set as B. For a two-tailed test, specify the alternative “`two.sided`”. If the alternative is not specified, then a default value of “`two.sided`” is assumed for `RSDT_power()`.

In the `RSDT_power()` example below, the arguments supplied are the expected case score on task A and task B (`case_a`, `case_b`), specifying an inter-task discrepancy of 2 standard deviations, for a sample size of 15, assuming an inter-task correlation of `.6` in the control sample. The alternative “`two.sided`” specifies a two-tailed test. The level of power is returned by the function (0.5049).


```
R> RSDT_power(case_a = -2,
+ case_b = 0,
+ r_ab = 0.6,
+ sample_size = 16,
+ alternative = "two.sided")
```

```
OUTPUT >
[1] 0.5049
```

`RSDT_power()` is simulation-based, and may take some time to complete. The number of simulations can be adjusted by changing `nsim` from its default value of 10000. Lower `nsim` values will produce less precise but faster estimates, and higher values will produce more precise but slower estimates. The power estimates for the RSDT provided in Table A2 were obtained with `nsim` set to 1000000.

Note that, because `RSDT_power()` is simulation-based, it is not possible to calculate the required sample size for a desired level of power using this function.

Appendix II: Power tables for the Test of Deficit and Revised Standardised Difference Test

Table A1 shows exact power for the Test of Deficit (TD: Crawford & Howell, 1998), for a range of likely parameters. Power for the TD is calculated from the formula:

$$power = 1 - \beta = T_{n-1} \left(t_{\alpha, n-1} \left| \frac{x^* - \bar{x}}{\sigma \sqrt{\frac{n+1}{n}}} \right| \right)$$

Where $T_{n-1}(\cdot|\theta)$ is the cumulative distribution function for the non-central t-distribution with $n - 1$ degrees of freedom and non-centrality parameter $TD = \frac{x^* - \bar{x}}{\sigma \sqrt{\frac{n+1}{n}}}$ is the α quantile of the *central* t-distribution on $n - 1$ degrees of freedom. See Harrison & Brady (2004) for discussion on using the non-central t-distribution for power calculations for t-tests.

Note that the Bayesian Test of Deficit (Crawford et al., 2011) produces equivalent results to the TD, and has frequentist properties, so may be considered to have the same effective levels of power as the TD. Note also that the Unstandardised Difference Test (UDT) (Crawford & Garthwaite, 2005) is identical to applying the TD to between-task discrepancy scores. Power for the UDT is thus identical with

power the TD, but related to the effect size of a between-task discrepancy (Z_{DCC}) rather than the effect size of a deficit (Z_{CC}).

Table A2 shows estimated power for the Revised Standardised Difference Test (RSDT: Crawford & Garthwaite, 2005), for a range of likely parameters. Power for the RSDT was estimated using Monte Carlo simulation, by generating $n+1$ random draws from a bivariate normal distribution, imposing the desired Z_{CC} on the $n+1$ th draw and comparing that observation to the n others using the RSDT test formula. This was repeated 1 000 000 times, and power was found by dividing the number of significant tests by the total number of tests run, for each parameter combination of interest. Code used for calculations and simulations is available at [github](#)⁶.

Table A1: Power for the Test of Deficit (Crawford & Howell, 1998).

Tails	α	Z_{CC}	Control sample size (n)										
			4	6	8	10	12	14	16	18	20	22	24
One	0.05	0.5	9.89	10.73	11.17	11.45	11.64	11.77	11.87	11.95	12.02	12.07	12.12
		1.0	17.43	20.04	21.43	22.30	22.88	23.31	23.63	23.88	24.09	24.25	24.39
		1.5	27.65	32.91	35.67	37.36	38.49	39.31	39.92	40.40	40.78	41.09	41.35
		2.0	39.91	48.12	52.22	54.64	56.23	57.36	58.20	58.84	59.36	59.78	60.13
		2.5	53.02	63.51	68.33	71.04	72.76	73.95	74.82	75.49	76.01	76.43	76.78
		3.0	65.60	76.88	81.50	83.92	85.40	86.38	87.09	87.61	88.02	88.35	88.62
		3.5	76.48	86.91	90.57	92.33	93.33	93.98	94.43	94.76	95.01	95.20	95.36
		4.0	85.02	93.41	95.83	96.87	97.43	97.77	97.99	98.16	98.28	98.37	98.44
	0.10	0.5	18.58	19.59	20.11	20.43	20.64	20.79	20.90	20.99	21.07	21.12	21.17
		1.0	30.54	33.27	34.66	35.50	36.06	36.46	36.76	37.00	37.19	37.34	37.47
		1.5	44.89	49.60	51.90	53.27	54.18	54.82	55.30	55.68	55.97	56.22	56.42
		2.0	59.80	65.94	68.77	70.39	71.44	72.18	72.73	73.15	73.48	73.75	73.98
		2.5	73.28	79.68	82.38	83.86	84.79	85.43	85.90	86.25	86.53	86.75	86.94
		3.0	83.91	89.40	91.46	92.52	93.17	93.59	93.90	94.13	94.31	94.45	94.57
		3.5	91.26	95.20	96.48	97.08	97.43	97.66	97.82	97.93	98.02	98.09	98.15
		4.0	95.73	98.12	98.77	99.05	99.20	99.29	99.36	99.40	99.44	99.46	99.48
Two	0.05	0.5	6.22	6.68	6.95	7.12	7.24	7.33	7.40	7.45	7.50	7.53	7.56
		1.0	9.90	11.84	12.97	13.70	14.21	14.59	14.87	15.10	15.28	15.43	15.56
		1.5	16.00	20.56	23.21	24.90	26.07	26.92	27.57	28.07	28.49	28.82	29.11
		2.0	24.31	32.48	37.07	39.92	41.85	43.24	44.28	45.09	45.75	46.28	46.72
		2.5	34.35	46.46	52.83	56.61	59.09	60.83	62.11	63.10	63.88	64.51	65.04
		3.0	45.36	60.77	68.08	72.14	74.68	76.40	77.64	78.58	79.31	79.89	80.37
		3.5	56.48	73.62	80.70	84.29	86.39	87.76	88.71	89.41	89.95	90.37	90.71
		4.0	66.89	83.80	89.64	92.28	93.71	94.59	95.18	95.59	95.90	96.14	96.33
	0.10	0.5	12.12	12.74	13.08	13.29	13.44	13.54	13.62	13.68	13.74	13.78	13.81
		1.0	18.30	20.73	22.05	22.87	23.42	23.83	24.14	24.38	24.57	24.73	24.86
		1.5	27.94	33.12	35.84	37.51	38.63	39.43	40.04	40.51	40.89	41.20	41.46
		2.0	39.99	48.17	52.25	54.67	56.26	57.38	58.22	58.86	59.38	59.80	60.15
		2.5	53.04	63.52	68.33	71.04	72.76	73.96	74.83	75.49	76.01	76.43	76.78
		3.0	65.60	76.89	81.50	83.92	85.40	86.38	87.09	87.61	88.02	88.35	88.62
		3.5	76.48	86.91	90.57	92.33	93.33	93.98	94.43	94.76	95.01	95.20	95.36
		4.0	85.02	93.41	95.83	96.87	97.43	97.77	97.99	98.16	98.28	98.37	98.44

Note: $Z_{CC} = \frac{x^* - \bar{x}}{s_x}$, i.e. the standardised case score for the task of interest.

Table A2: Power for the Revised Standardised Difference Test (Crawford & Garthwaite, 2005).

Tails	α	Z_{DCC}	Control sample size (n)										
			4	6	8	10	12	14	16	18	20	22	24
One	0.05	0.5	10.43	10.47	10.83	11.16	11.40	11.59	11.71	11.81	11.85	11.90	12.02
		1.0	16.78	18.84	20.42	21.50	22.20	22.77	23.11	23.39	23.65	23.93	24.01
		1.5	25.14	30.33	33.57	35.60	37.13	38.07	38.91	39.43	39.89	40.31	40.64
		2.0	34.86	43.72	48.80	51.95	54.00	55.38	56.46	57.37	58.03	58.49	59.01
		2.5	45.26	57.49	63.87	67.35	69.77	71.42	72.60	73.61	74.28	74.84	75.31
		3.0	55.49	69.84	76.47	80.10	82.27	83.71	84.80	85.59	86.19	86.82	87.19
		3.5	64.84	79.79	85.78	88.85	90.52	91.69	92.51	93.11	93.56	93.90	94.18
		4.0	72.89	87.03	91.87	94.16	95.43	96.14	96.68	97.08	97.33	97.50	97.70
	0.10	0.5	19.01	19.22	19.66	20.02	20.37	20.48	20.61	20.78	20.91	20.98	21.02
		1.0	29.52	31.78	33.36	34.52	35.17	35.71	36.19	36.55	36.70	36.91	37.08
		1.5	41.63	46.72	49.56	51.40	52.60	53.51	54.15	54.60	55.06	55.38	55.66
		2.0	54.39	61.64	65.48	67.73	69.26	70.31	71.22	71.76	72.10	72.56	72.91
		2.5	65.99	74.54	78.57	80.80	82.30	83.42	84.10	84.63	85.16	85.47	85.72
		3.0	75.60	84.14	87.76	89.71	90.90	91.73	92.33	92.73	93.09	93.34	93.58
		3.5	82.94	90.64	93.49	94.89	95.75	96.28	96.66	96.96	97.16	97.33	97.50
		4.0	88.41	94.64	96.68	97.60	98.14	98.47	98.70	98.86	98.96	99.06	99.12
Two	0.05	0.5	7.07	6.67	6.85	7.03	7.15	7.23	7.33	7.38	7.43	7.45	7.49
		1.0	9.85	11.01	12.31	13.16	13.70	14.14	14.59	14.80	15.02	15.24	15.36
		1.5	14.41	18.47	21.47	23.49	24.88	25.99	26.72	27.37	27.84	28.22	28.56
		2.0	20.36	28.50	33.99	37.50	39.76	41.53	42.77	43.73	44.55	45.07	45.69
		2.5	27.50	40.40	48.14	52.87	56.06	58.27	59.81	61.10	62.03	62.90	63.52
		3.0	35.28	52.75	62.09	67.48	71.06	73.19	74.86	76.13	77.13	77.93	78.64
		3.5	43.45	64.41	74.24	79.44	82.57	84.58	86.03	87.14	87.93	88.52	89.03
		4.0	51.48	74.39	83.56	87.91	90.42	91.96	93.01	93.77	94.34	94.72	95.10
	0.10	0.5	13.59	12.80	12.99	13.13	13.33	13.40	13.50	13.60	13.67	13.72	13.69
		1.0	18.57	19.88	21.24	22.18	22.83	23.33	23.68	23.94	24.23	24.33	24.62
		1.5	26.16	30.71	33.94	35.95	37.28	38.33	39.00	39.66	40.11	40.43	40.86
		2.0	35.49	44.01	49.00	52.00	54.02	55.46	56.61	57.41	58.13	58.67	59.09
		2.5	45.74	57.70	63.84	67.50	69.88	71.49	72.65	73.50	74.32	74.88	75.30
		3.0	55.78	69.87	76.48	80.07	82.20	83.75	84.80	85.54	86.17	86.71	87.14
		3.5	65.07	79.80	85.77	88.82	90.54	91.68	92.47	93.07	93.58	93.95	94.22
		4.0	73.06	87.13	91.92	94.15	95.38	96.13	96.68	97.03	97.32	97.53	97.69

Note: $Z_{DCC} = (Z_{CC_Y} - Z_{CC_X})/\sqrt{2 - 2r}$, i.e. the standardised case score for the intertask discrepancy of interest where z_x^* and z_y^* are the standardised case scores for the tasks of interest. Higher power for small samples and small effect sizes (see two-tailed tests when $n = 4$ and $Z_{DCC} = 0.5$) is due to a loss of control of Type I errors (Crawford & Garthwaite, 2005). Power calculated with 1 000 000 Monte Carlo simulations.