Rounded Input Variables, Exact Test Statistics (RIVETS): A Technique For Detecting Hand-Calculated Results in Published Research.

Nicholas J. L. Brown, University of Groningen

James A. J. Heathers, Northeastern University

Abstract

Most statistical software packages report the input values to statistical tests (e.g., means and standard deviations for an unpaired t test) in a rounded form (e.g., to two decimal places), with this rounding having been performed after the test statistic has been calculated.

However, in some cases, the input values are rounded before the test statistic is calculated, most likely because of some form of manual intervention by the researcher. We describe a method that enables the probabilistic identification of detecting this rounding, the conditions required for this method to be applicable, the tests where pre-calculation rounding can be detected, and the implications of its detection.

Consider the following (imaginary) fragment from a published article in *Airborne Farm Animal Review*, describing the effects of different forms of flying on pigs:

Pigs in the self-piloting condition ($N = 55$) ate more kilograms of feed ($M = 4.53$, $SD = 1.07$) in the 24-hour period after landing than did pigs in the drone-transported condition ($N = 53$, $M = 3.73$, $SD = 0.92$), $t(106) = 4.16$, $p < .001$.

This calculation is precisely correct, and can be verified by any software package that can perform an independent samples t test. If one inputs the exact numbers above for the mean (M), standard deviation (SD), and sample size (N), the computed t statistic will be 4.159401, which rounds to 4.16.

However, a degree of hidden imprecision is often involved in such calculations. Because the input values for the means and SD s are typically rounded when reported, the actual values that contributed to the calculation of the t statistic may have been larger or smaller (e.g., the mean amount of feed consumed by the self-piloting pigs, reported as 4.53 when rounded to two decimal places, might have been more specifically reported as 4.527 or 4.5273 or 4.52728 with different rounding). Crucially, this loss of precision means that a range of values for the t statistic can be consistent with the same set of rounded means and SD s.

In this case, the maximum possible value for the test statistic is calculated from the maximum possible difference between means combined with the minimum possible SD s, which (when rounded to six decimal places, although in the general case the fractional part could be arbitrarily long¹) are ($M = 4.534999^2$, $SD = 1.065001$) in the first condition and ($M = 3.725001$, $SD = 0.915001$) in the second³. These values give the same numbers as

¹ For sample sizes below several thousand and rounding to two decimal places, adding yet more precision—in the form of extending the number that is added or subtracted to the mean value with 9s beyond the sixth decimal place—does not change the minimum or maximum t statistic to any appreciable degree.

² We add or subtract 0.004999 rather than .005000 in our simulation code to avoid issues caused by the rounding protocols of different software packages. For example, both SPSS and R round 0.275 to 0.28, but SPSS rounds 0.265 to 0.27, whereas R rounds 0.265 to 0.26.

³ Adding 0.004999 to the (rounded) larger mean and subtracting 0.004999 from the smaller mean gives the greatest possible increase (almost 0.01) in the difference between the means, which appears

reported in the article after rounding to two decimal places, but would result in a rounded t statistic of 4.23⁴. A similar calculation with the opposite adjustments (moving the means closer together and increasing the SDs) shows that the minimum possible rounded t statistic is 4.09. These calculations are summarized in Table 1.

Table 1. Demonstration of the derivation of the minimum and maximum possible t statistics corresponding to the reported comparisons of the post-flight feed consumption of the pigs.

Case	Mean 1	SD 1	Mean 2	SD 2	t (exact)	t (rounded)
Minimum t	4.525001	1.074999	3.734999	0.924999	4.087029	4.09
Reported t	4.530000	1.070000	3.730000	0.920000	4.159401	4.16
Maximum t	4.534999	1.065001	3.725001	0.925001	4.232497	4.23

We calculated one million t statistics using uniformly distributed random means and SDs identical to those reported in the article (to two decimal places), and found that for this particular combination of $M/SD/N$ values only 16.50% of cases produced a rounded t statistic of 4.16. This percentage varies considerably with different input values, and we have not been able to identify a good rule of thumb to estimate it in the general case. It tends to decrease as the difference between the means increases, to increase as the SDs increase, and to decrease as the sample sizes increase. It does not appear to depend on the number of reported decimal places.

in the numerator of the formula for the t statistic. Similarly, subtracting 0.004999 from both standard deviations gives the greatest possible reduction in the pooled standard error term, which appears in the denominator. Hence, this combination of operations maximizes the t statistic without affecting the rounded (to two decimal places) mean and SD values. For the inverse operation, one would subtract 0.004999 from the (rounded) larger mean and add 0.004999 to the smaller mean, thus decreasing the difference between the means by almost 0.01, and add 0.004999 to both standard deviations.

⁴ All of the code needed to reproduce the analyses in the present article can be found at <https://osf.io/ubwm3>.

This observation provides us with a heuristic for identifying sets of results where the test statistic may have been calculated “by hand”—for whatever reason—from the rounded descriptive statistics, versus the more normal process of allowing a statistical package to calculate everything in one run from the raw data points. For example, imagine that the numbers that appear in the extract cited earlier were just the first of a set of reported results from a table in the same article, describing a series of observed behaviors among the flying pigs (besides the amount of feed eaten, these were the numbers of hours during which the pigs were observed to sleep, grunt, and dig during the next day), which we reproduce here as our Table 2:

Table 2. Measures of behavior of flying pigs after landing, by aerial navigation condition.

Variable (units)	Self-piloting ($N = 55$)		Drone-transported ($N = 53$)		Difference		
	M	SD	M	SD	$t(106)$	p	$p(\text{ident})$
Feed (kg)	4.53	1.07	3.73	0.92	4.16	<0.0001	.1657
Sleeping (h)	8.18	2.12	9.57	1.89	-3.59	0.0005	.3258
Grunting (h)	1.76	0.29	1.38	0.34	6.42	<0.0001	.0478
Digging (h)	1.47	0.25	1.32	0.33	2.89	0.0088	.0514

In Table 2, the last column, designated $p(\text{ident})$, represents the fraction of t statistics in our simulations that were identical to the ones reported in the table for the given M and SD values; put another way, it represents the chance that any given run of our simulation would return an identical t statistic. We believe that it is appropriate to claim that the chance that all four t statistics would be identical to the ones calculated, assuming that the means and SDs were rounded after the calculation was performed, is the product of these four fractions (i.e., $0.1657 \times 0.3258 \times 0.0480 \times 0.0514$), which is 0.00013. That is, we feel confident in rejecting the null hypothesis that all four of these t statistics are the products of a standard analysis procedure in which the authors fed their raw data (the observed values of Feed, Sleeping, Grunting, and Digging for each pig) into a software package that then simultaneously

generated the means, SDs, and test statistics, and rounded these values only after the calculation was complete. Our alternative hypothesis is that the means and SDs were written down as exact quantities (rounded to two decimal places) at some point, and those numbers were then used as the basis for calculating the t statistics. We refer to the test of these hypotheses as RIVETS (Rounded Input Variables, Exact Test Statistics), for want of any better ideas for a name.

Other statistics

Use of the RIVETS test is not confined to the outcomes of independent samples t tests. The same principle can be used to check many other types of test statistic where truncated input values are reported. In such cases, an omnibus p value for RIVETS⁵ can be calculated from $p(\text{ident})$ values using the same types of simulation as we described above. Some examples are:

- Regressions (or other linear models) where statistical significance is evaluated using a t test (for which the regression coefficient is divided by its standard error); see our example below.
- Between-subjects one-way or two-way ANOVA, if the results presented can be recalculated from the summary statistics. For instance, if three groups (all with $N = 20$ and $SD = 5.26$) with means of 10.41, 8.41, and 6.41 are subjected to a one-way ANOVA, recalculation with the full range of possible values for the means and SDs gives a range for the resulting F statistic of 2.87 to 2.91. In our accompanying code we provide an example function (named `f_range`) that uses the R package `rpsychi` to calculate such ranges for one- and two-way ANOVAs, subject to certain limitations.
- In some cases, Pearson correlations rounded to two decimal places are compatible with multiple p values. For example, $r(50) = .07$ has a p value of .62, but that reported r value could range from .0651 ($p = .65$) to .0749 ($p = .60$). However, our informal

⁵ Note that the p value of the original statistic (e.g., the p value associated with the t tests in our initial example) is typically irrelevant, unless that p value is what takes the role of the “test statistic” in our terms (as is the case with Pearson correlations, for example).

calculations show that when the p value is below about .02 it tends to be stable regardless of the addition or subtraction of .0049 to the correlation coefficient.

- In a few cases, $p(\text{ident})$ can be calculated from chi-squared statistics if the inputs to the chi-squared test are reported as truncated percentages. For example, consider a sample of 1,847 people for whom it is reported that “52% said they would vote for Alice and 48% said they would vote for Bob, $X^2(1) = 2.89$, $p = .09$,” in a context that suggests that the percentages have been rounded to the nearest integer. In this case, the vote tallies could range from (Alice = 970, Bob = 877, $X^2(1) = 4.68$, $p = .03$) through (Alice = 951, Bob = 896, $X^2(1) = 1.64$, $p = .20$) while still matching the rounded integer percentages.

Granularity considerations

An important concern when using the ideas behind RIVETS is that the means of input variables that are not measured continuously (e.g., Likert-type integer data in psychology) may not be able to take on an arbitrarily large number of distinct values (cf. Brown & Heathers, 2017). With a one-item Likert-type variable, a sample size of less than 100, and reporting to two decimal places, there will typically only be only one possible unrounded mean value that leads to the reported rounded mean; for example, with $N = 41$ and a rounded mean of 2.88, the only possible unrounded mean is $\overline{2.87804}$, corresponding to a total score on this item for all 41 participants of 118.

However, with larger samples, multiple Likert-type items, or suitably precise measurement of continuous quantities, this number increases considerably. In our flying pigs example, we assumed that our quantities were measured with quite high precision (e.g., the amount of feed consumed could have been weighed to the nearest gram, and the time spent sleeping could have been measured to the nearest minute), so that the number of distinct input values would be quite large. For example, the measure of feed in the self-piloting condition, with a precision of 1g and $N = 55$ pigs, allows for 548 unique means that round to 4.53, while the mean of 9.57 hours in the drone-transported condition with a precision of 1 minute and $N = 53$ pigs allows for 349 unique means that round to 9.67. More generally, with k Likert-type items, a sample size of N , and rounding to d decimal places, the minimum

number of unique unrounded mean values that can give the reported rounded value is the integer part of $(k \times N) \div 10^d$, minus 1 or 2 that are lost to rounding at the edges.

When the number of possible unrounded input values is restricted, the true $p(\text{ident})$ value may be somewhat higher or lower than what would be estimated if the simulation code were to assume that the input values are essentially continuous. (In the pathological case where there is only one possible unrounded value for each of the inputs, the true value of $p(\text{ident})$ is equal to 1.) Our simulations of the independent samples t test for the flying pigs showed that when the number of distinct possible unrounded means was forced to 2 or 3 in each condition, $p(\text{ident})$ was estimated at around .42 or .14 respectively (versus .1657 in Table 2). However, once the number of distinct unrounded means exceeded about 12—which, assuming rounding to two decimal places, corresponds to an N of approximately 1,300 with single-item Likert-type data, 250 with a five-item Likert scale, or 60 with a precision of 20g when measuring a mean feed amount of the order of a kilogram—the results did not differ substantially from the case where the means were treated as continuous quantities, which can be considerably easier to program.

If the minimum sample/item count sizes are not met (or, in the case of different statistical tests, there appears to be another issue with low granularity), a researcher who wishes to perform RIVETS simulations will need to consider sampling from an exhaustive list of possible input values, rather than simply generating continuous random numbers within the appropriate range. This can be a challenging task in the case of standard deviations (cf. Anaya, 2016; Heathers, Anaya, van der Zee, & Brown, 2018; Wilner, Wood, & Simons, 2018); happily, however, the number of unrounded standard deviation values that can give any reported rounded value is typically considerably higher than the number of means, so that most sample sizes that are likely to be encountered in published research will provide the required number of distinct SDs for the assumption of their equivalence to continuous values to be met. In some other cases (e.g., regression coefficients), it may be impossible to generate the full set of possible input values for a given variable. In any case, researchers should to satisfy themselves (and their readers) that the number of possible values in the case at hand is sufficiently large for the use of continuous values to be justified.

In summary, any RIVETS-based analysis of published results should include a principled consideration of the possible effects of the granularity of the input variables, and either sample only from possible values, or demonstrate that the number of possible input

values is sufficiently diverse (we suggest at least 20 to be on the safe side) to allow continuous random numbers to be used. This will generally mean that the simulation code for any given RIVETS analysis will need to be custom-built for each case. For this reason, the example code that accompanies the present article is supplied purely on an “as-is” basis, and not for example as an R package. We suggest that researchers who report analyses based on RIVETS that use simulation to estimate $p(\text{ident})$ should include a statement about how they took granularity into account. We provide an example of such a statement in our example from the published literature below.

Possible causes of significant RIVETS tests

There are several reasons why a researcher might truncate data and consequently produce a set of descriptive and test statistics that will be flagged by RIVETS. They might choose to calculate their test statistics manually due to some unfortunate combination of circumstances, perhaps involving the loss of the raw data, with only some hastily printed descriptive statistics being salvageable. A miscommunication between workgroups or collaborators (perhaps between a primary researcher and an analyst) might mean that descriptive statistics are available while raw data is not; since the accuracy of the calculations is typically not going to be severely affected, the analyst may proceed anyway for reasons of speed or convenience. Such a situation could also arise if data cannot be transferred between all of the members of a research collaboration. For example, the raw data might be so sensitive that only members of lab X are ethically mandated to inspect them, so an investigator from lab X might provide truncated summary values to the overall project analyst in lab Y who then calculates the test statistics.

A further possibility is that the process of rounding before calculating the means is a feature of the software being used, although it would be very unusual and such a discarding of precision would most likely indicate a programming oversight. Alternatively, some older researchers might prefer to work out their t statistics using a calculator, or a table of square roots, because they have always done so. However, both of these scenarios seem to us to be somewhat unlikely. All modern statistical analysis packages with which we are familiar make it very easy to report means, SDs, and t statistics (or in a single step directly from raw data items, and we believe that almost all researchers have been using such software more or less exclusively to calculate their results for at least the last 20 years. Certainly we have never

seen a method section that states that “Preliminary analyses were performed with software package X, but due to the limitations of this package, the summary statistics were rounded and used to calculate the test statistics by hand.”

There are, of course, more problematic possibilities. For instance, a researcher might use a “quick and dirty” local calculation method (e.g., a t test calculation formula in an Excel spreadsheet) or an online calculator, with input values truncated for convenience, to “nudge” a crucial statistic over a threshold for significance (say, from $p = .07$ to $p = .04$) by iterating through a series of jittered values. Of course, by itself a positive RIVETS test should not be regarded as evidence of intent to deceive; rather, it simply suggests that truncation has occurred at some point. However, we have yet to see a method section that reports something to the effect of “at this point, results were truncated and then hand-calculated” after describing an earlier analytical process within a statistical package. Thus, a low $p(\text{ident})$ value for the RIVETS test will, in most cases, imply at least a deviation from a reported analysis plan.

Example: Stewart (2003)

As a practical example of the use of RIVETS, we examine an article by Stewart (2003), who reported, in his Tables 4 and 5, a total of 36 pairs of regression coefficients (b) and standard errors (SE), together with their associated t statistics (defined as the ratio b/SE). As the spreadsheet that we have made available in our OSF repository shows, in every case the reported t statistic is exactly what would be obtained if the reported b and SE values had been rounded to three decimal places before the t statistic was calculated, rather than having been rounded after the ratio between them had been calculated by a statistical software package. We performed some simulations (one million pairs of b and SE values based on the values reported in Stewart’s Tables 4 and 5, with uniformly distributed random disturbance terms⁶ in the asymptotic range $[-0.0005, 0.0005]$), and obtained the fraction of exactly

⁶ We believe that the b and SE variables can be treated as continuous in this case (cf. our section entitled “Granularity considerations”) because the sample sizes are very large (10,578 students nested within 528 schools), and the hierarchical linear models reported by Stewart (2003) include at least 14 predictor variables. It therefore seems reasonable to expect that there would be a large number of numerically possible b and SE values that could round to the reported values shown in Table 3.

matching values shown in the right-hand column of our Table 3. It can be seen that in general, the individual fractions are smaller (corresponding to a lower likelihood in each case that the rounding was done after calculation of the t statistic) than the numbers that we obtained when simulating independent t tests in the previous example. There are some repetitions in Stewart's Tables 4 and 5, with the "Within Schools" section being identical in each, but there are still 22 unique variables. Multiplying the $p(\text{ident})$ values for these 22 variables together gives an overall p value associated with the test of the null hypothesis that the 22 corresponding pairs of input values (b and SE) were rounded after calculating the corresponding t statistic, rather than having been written down at some point with only three decimal places of precision before the t statistic was calculated, of around 3.5×10^{-44} .

Table 3. Variable labels, regression coefficients, standard errors, and t statistics from Stewart's (2003) Tables 4 and 5, together with the fraction of cases in which our simulations gave exactly the same t statistic, after rounding to three decimal places.

Table	Variable	b	SE	t (calc)	p (ident)
4	Grand mean	0.941	0.019	49.526	.0004
4, 5	School attachment	-0.212	0.031	-6.839	.0045
4, 5	School commitment	-0.145	0.042	-3.452	.0120
4, 5	Belief in school rules	-0.397	0.039	-10.179	.0039
4, 5	School involvement	-0.011	0.014	-0.786	.0138
4, 5	Positive peers	-0.431	0.039	-11.051	.0035
4, 5	Parental school involvement	-0.138	0.043	-3.209	.0133
4, 5	Student's GPA	-0.108	0.047	-2.298	.0205
4, 5	Family structure (0 = single)	0.032	0.029	1.103	.0263
4, 5	Family income	-0.102	0.049	-2.082	.0239
4, 5	Gender (1 = male)	0.345	0.048	7.188	.0066
4, 5	African American (1 = AA)	0.099	0.049	2.020	.0242
4, 5	Asian (1 = Asian)	-0.103	0.049	-2.102	.0233
4, 5	Hispanic (1 = Hispanic)	0.034	0.029	1.172	.0248
4, 5	Native American (1 = NA)	0.013	0.016	0.813	.0160
5	Grand mean	0.928	0.021	44.190	.0005
5	Proportion nonwhite	0.039	0.041	0.951	.0410
5	School size	0.197	0.042	4.690	.0090
5	School poverty	0.017	0.026	0.654	.0261
5	School location (1 = urban)	0.173	0.044	3.932	.0112
5	School social problems	0.032	0.029	1.103	.0263
5	School cohesion	-0.035	0.029	-1.207	.0239
5	School attachment	-0.212	0.031	-6.839	.0045

Notes

Table: Table number from Stewart (2003).

b : column headed "Coefficient" in Stewart's (2003) tables.

t (calc): The calculated ratio obtained when dividing the reported b value by the reported SE value. In every case, this is identical to the number reported in Stewart's (2003) table.

p (ident): Fraction of t statistics from our simulations that matched Stewart's (2003) original values after rounding to three decimal places.

Conclusion

Evidence from error or anomaly detection tests is often probabilistic in nature (as is, indeed, most data in the social sciences, where estimation of error is typically one of the main aims of a study). A positive RIVETS test—defined as a low $p(\text{ident})$ value—cannot, by itself, tell us anything definitive about what may have led to it. However, when such positive tests are encountered, it may be worthwhile to contact the author of the article in question and ask for details of the method used, and perhaps also a copy of the data set so that the method can be used to reproduce the published results. We hope that RIVETS, like the other tools that we have been involved in developing (Brown & Heathers, 2017; Heathers et al., 2018; Heathers & Brown, 2019) may prove useful in explaining apparently anomalous results in published research in a variety of fields.

References

- Anaya, J. (2016). The GRIMMER test: A method for testing the validity of reported measures of variability. *PeerJ Preprints*, 4, e2400v1.
<https://doi.org/10.7287/peerj.preprints.2400v1.6>
- Brown, N. J. L., & Heathers, J. A. J. (2017). The GRIM test: A simple technique detects numerous anomalies in the reporting of results in psychology. *Social Psychological and Personality Science*, 8, 363–369.
<https://doi.org/10.1371/10.1177/1948550616673876>
- Heathers, J. A. J., Anaya, J., van der Zee, T., & Brown, N. J. L. (2018). Recovering data from summary statistics: Sample Parameter Reconstruction via Iterative TEchniques (SPRITE). *PeerJ Preprints*, 6, e26968v1.
<https://doi.org/10.7287/peerj.preprints.26968v1>
- Heathers, J. A. J., & Brown, N. J. L. (2019). *DEBIT: A simple consistency test for binary data*. Retrieved from <http://osf.io/pm825>
- Stewart, E. A. (2003). School social bonds, school climate, and school misbehavior: A multilevel analysis. *Justice Quarterly*, 20, 575–604.
<https://doi.org/10.1080/07418820300095621>
- Wilner, S., Wood, K., & Simons, D. J. (2018, January 29). Complete recovery of values in Diophantine systems (CORVIDS). *PsyArXiv*. <https://doi.org/10.31234/osf.io/7shr8>