The GRIM test: A simple technique detects numerous anomalies in the reporting of results in psychology

Nicholas J. L. Brown (*)
University Medical Center Groningen

James A. J. Heathers
Division of Cardiology and Intensive Therapy, Poznań University of Medical Sciences
University of Sydney
(*) Corresponding author: E-mail: nick.brown@free.fr


#### Abstract

We present a simple mathematical technique that we call GRIM (Granularity-Related Inconsistency of Means) for verifying the summary statistics of published research reports in psychology. This technique evaluates whether the reported means of integer data such as Likerttype scales are consistent with the given sample size and number of items. We tested this technique with a sample of 260 recent articles in leading journals within empirical psychology. Of the subset of articles that were amenable to testing with the GRIM technique ( $N=71$ ), around half $(N=36 ; 50.7 \%)$ appeared to contain at least one reported mean inconsistent with the reported sample sizes and scale characteristics, and more than $20 \%(N=16)$ contained multiple such inconsistencies. We requested the data sets corresponding to 21 of these articles, receiving positive responses in 9 cases. We were able to confirm the presence of at least one reporting error in all cases, with 2 articles requiring extensive corrections. The implications for the reliability and replicability of empirical psychology are discussed.


Consider the following extract from a recent article in the Journal of Porcine Aviation Potential. The authors' principal hypothesis was that drinking Kool-Aid ${ }^{1}$ increases people's willingness to believe that pigs can fly.

Fifty-five undergraduates took part in this study in return for course credit. Participants were randomly assigned to drink eight ounces of water, which either contained (experimental condition, $N=28$ ) or did not contain (control condition, $N=27$ ) 17 g of cherry flavor Kool-Aid powder. Fifteen minutes after consuming the beverage, participants responded to the question, "To what extent do you believe that pigs can fly?" on a 7-point Likert-type scale from 1 (Not at all likely) to 7 (Overwhelmingly likely). Participants in the "drunk the Kool-Aid" condition reported a significantly stronger belief in the ability of pigs to fly $(M=5.19, S D=1.34)$ than those in the control condition $(M=3.87, S D=1.41), t(53)=3.56, p<.001$.

These results (and similarly improbable but provocative data in real articles) may garner both intense public interest and skepticism. They also sometimes provoke speculation over their fidelity; this speculation frequently continues in public fora such as PubPeer and Twitter in the absence of re-analysis of the data in question. However, in a subset of cases, it is possible to determine the fidelity of scale data directly from the descriptive statistics.

The case cited above, for example, seems superficially reasonable but actually describes a situation which is mathematically impossible. The reported means represent either errors of transcription, some version of misreporting, or the deliberate manipulation of results.
${ }^{1}$ The popular metaphor referred to in this example may be based on a distortion of reality (Reiterman \& Jacobs, 1982). We apologize to the good people at Kraft Foods, manufacturers of Kool-Aid, for repeating this myth one more time here to make a rhetorical point.

Specifically, the mean of the 28 participants in the experimental condition, reported as 5.19 , cannot be correct. Since all responses were integers between 1 and 7 , the total of the response scores across all participants must also be an integer in the range $28-196$. The two integers that give a result closest to the reported mean of 5.19 (which will typically have been subjected to rounding) are 145 and 146. However, 145 divided by 28 is 5.17857142 , which conventional rounding returns as 5.18. Likewise, 146 divided by 28 is $5.21 \overline{428571}$, which rounds to 5.21 . That is, there is no combination of responses to the question that can give a mean of 5.19 when correctly rounded. Similar considerations apply to the reported mean of 3.87 in the control condition. Multiplying this value by the sample size (27) gives 104.49 , suggesting that the total score across participants must have been either 104 or 105 . But 104 divided by 27 is $3 . \overline{851}$, which rounds to 3.85 , and 105 divided by 27 is $3.88 \overline{8}$, which rounds to 3.89 .

In this article, we will use the term inconsistent to refer to reported means of integer data whose value, appropriately rounded, cannot be reconciled with the stated sample size. We first introduce the general background to and calculation of what we term the Granularity-Related Inconsistent Means (GRIM) test. Next, we report on the results of an analysis using the GRIM test of a number of published articles from leading psychological journals. Finally, we discuss the implications of these results for the published literature in empirical psychology.

## General description of the GRIM approach for reanalyzing published data

The crux of this method lies in the transition from ordinal to continuous data. Scale data collected in psychology typically lie on an ordinal scale - that is, the recorded values are in rank order but are arbitrary, such that the number corresponding to any scale item has no technical significance beyond its ability to establish a position on a continum relative to the other
numbers. For example, if we use a typical Likert-type scale, and grade an opinion on a chosen subject from 1 (strongly disagree) to 5 (strongly agree), the difference between 2 (disagree) and 3 (neither agree nor disagree) is not directly equivalent to the difference between 1 and 2 .

While the limits of ordinal data in measurement have been extensively discussed in measurement theory (e.g., Coombs, 1960; Thurstone, 1927), this discussion is largely separate from research practice. It is presently common to a) treat ordinal scale measures as continuous variables, b) calculate means and standard deviations from these numbers, and c) subject those values to null-hypothesis significance testing. Although discussions as to the general validity of such scales, whether composed of single or multiple items, appear in the literature from time to time (e.g., Jamieson, 2004; Carifio \& Perla, 2007), the position of these measures as the dominant paradigm in psychometric testing does not seem to be under any immediate threat.

One often-overlooked property of data derived from Likert-type scales is their granularity-that is, the numerical separation between possible values of the summary statistics. Here, we consider the example of the mean. As our initial example demonstrates, within typical ordinal scale data the smallest unit by which two means can differ is an inverse function of the number of participants and the number of scale items. In other words, if we consider a Likerttype scale administered to 10 people averaged across three items, the smallest amount by which two mean scores can differ (the granularity) is $1 /(10 \times 3)=0.03 \overline{3}$. If means are reported to two decimal places (as is the case in the great majority of psychology journals), then, although numerically there are 100 possible two-digit fractional parts of a mean $M$ in the range $1 \leq M<2$ ( $1.00,1.01,1.02,1.03$, etc., up to 1.99 ), the possible outcomes from the process of division (to obtain the mean) and subsequent rounding to two decimal places are considerably fewer (1.00, $1.03,1.07,1.10,1.13$, etc., up to 1.97). In the typical case where means are reported to two
decimal places, if the number of participants $(N)$ is less than 100 and the measured quantity is an integer, then not all of the possible sequences of two digits can occur after the decimal point in correctly rounded fractions. (More generally, if the number of decimal places reported is $D$, then some combinations of digits will not be consistent if $N$ is less than $10^{D}$.)

This relation is always true for integer data that are recorded as single items, such as participants' ages in whole years, or a one-item measure of an attitude to a specific issue, as is frequently found in a manipulation check. When a composite measure is used, such as a scale with three Likert-type items on a scale of $1-7$ where the mean of the item scores is taken as the value of the measure, this mean value will not necessarily be an integer; instead, it will be some multiple of $(1 / L)$, where $L$ is the number of items in the measure. However, the range of possible values that this mean can take is still constrained (in the example, to $1.00,1.3 \overline{3}, 1.6 \overline{6}$, $2.00,2.3 \overline{3}$, etc.,) and so for any given sample size, the range of possible values for the mean of all participants is also constrained. For example, with a sample size of 15 and $L=3$, possible values for the mean are 1.00, 1.02 [rounded from $1.016 \overline{6}]$ ] 1.03 [rounded from $1.033 \overline{3}], 1.05$, $1.07,1.08,1.10$, etc. More generally, the range of means for a measure with $L$ items and a sample size of $N$ is identical to the range of means for a measure with one item and a sample size of $L \times N$. Thus, by multiplying the sample size by the number of items in the scale, composite measures can be analyzed using the GRIM technique in the same way as single items, although as the number of scale items increases, the maximum sample size for which this analysis is possible is correspondingly reduced, as the granularity decreases towards 0.01 . For example, a five-item measure with 25 participants has the same granularity ( 0.008 ) as a one-item measure with 125 participants, and hence scores on this measure are not typically GRIM-testable.

Note that the adjustment for sample size just mentioned is only required when the composite measure is presented as a mean. If it is presented as a total (e.g., a three-item measure on a scale of $1-7$ reported as a total in the range 3-21), no adjustment is necessary. Indeed, one of the attractions of the GRIM technique is that it is independent of the possible values that the measured variable can take, provided either that these are integers, or (in the case of composite items) that they can be represented as integers divided by a sufficiently small number of items. In particular, the number of possible responses to a single Likert-type item (such as five or seven) is irrelevant. This technique can also be applied to real or quasi-real values, such as ages in years, provided that these were measured as integers (or simple fractions thereof, with an adjustment to the sample size as discussed above) $)^{2}$. We refer to variables that are amenable to testing for inconsistencies by our technique as "GRIM-testable data."

[^0]
## Insert Figure 1 around here

Figure 1 shows the distribution of consistent (shown in white) and inconsistent (shown in black) means as a function of the sample size. Note that it only the 2-digit fractional component is linked to consistency; the integer part of the mean plays no role. The overall pattern is clear: As the sample size increases towards 100 , the number of fractional means that is consistent with that sample size also increases, and so the chance that any incorrectly-reported (due to a random error) mean will appear as an anomaly is reduced. However, even with quite large sample sizes, it is still possible to detect inconsistent means if an article contains multiple inconsistencies. For example, consider an article with $N=75$ and six reported means that have, in fact, been mistyped. For any one mean, there is a $75 \%$ chance that it will be consistent, but there is only a $17.8 \%$ chance that all six means will be consistent.

Our general formula, then, is that when the number of participants $(N)$ is multiplied by the number of items composing a measured quantity ( $L$, commonly equal to 1 ), and the means that are based on $N$ are reported to $D$ decimal places, then if $(L \times N)<10^{D}$, there exists some number of decimal fractions of length $D$ that cannot occur if the means are reported correctly. The number of inconsistent values is generally equal to $\left(10^{D}-N\right)$, although there are a few values of $N$ where this breaks down if means that end in 5 at the $(D+1)$ th place are allowed to be rounded either up or down (e.g., with $D=2$ and $N=88$, a total score of 297 gives a mean of exactly 3.125 . Using the most common rounding rule this terminal 5 would be rounded up to
3.13, but other rules exist which would see it rounded down to 3.12: indeed the round () function in the R programming language would give this result, since it rounds to the nearest even number in the last digit position). In all of the analyses reported in the present article, we conservatively allowed numbers ending in exactly 5 at the third decimal place to be rounded either up or down without treating the resulting reported means as inconsistent.

## A numerical demonstration

For readers who prefer to follow a worked example, we present here a simple method for performing the GRIM test to check the consistency of a mean. We assume that some quantity has been measured as an integer across a sample of participants and reported as a mean to two decimal places. For example:

$$
\text { Participants }(N=52) \text { responded to the manipulation check question, "To what extent did }
$$ you believe our story about the dog having eaten our homework?" on a 1-7 Likert-type scale. Results showed that they found our story convincing ( $M=6.28, S D=1.22$ ).

The first step is to take the mean and multiply it by the sample size. In this example, that gives $(6.28 \times 52)=326.56$. Next, round that product to the nearest integer (here, we round up to 327). Now, divide that integer by the sample size, rounding the result to two decimal places, giving $(327 / 52)=6.29$. Finally, compare this result with the original mean. If they are identical, then the mean is consistent with the sample size and integer data; if they are different, as in this case (6.28 versus 6.29 ), the mean is inconsistent.

When the quantity being measured is a composite Likert-type measure, or some other simple fraction, it may still be GRIM-testable. For example:

> Participants $(N=21)$ responded to three Likert-type items $(0=$ not at all, $4=$ extremely $)$ asking them how rich, famous, and successful they felt. These items were averaged into a single measure of fabulousness $(M=3.77, S D=0.63)$.

In this case, the measured quantity (the mean score for fabulousness) can take on the values 1.00 , $1.3 \overline{3}, 1.6 \overline{6}, 2.00,2.3 \overline{3}, 2.6 \overline{6}, 3.00$, etc. The granularity of this quantity is thus finer than if it had been reported as an integer (e.g., if the mean of the total scores for the three components, rather than the mean of the means of the three components, had been reported). However, the sample size is sufficiently small that we can still perform a GRIM test, by multiplying the sample size by the number of items that were averaged to make the composite measure (i.e., three) ${ }^{3}$ before performing the steps just indicated. Thus, in this case, we multiply the sample size of 21 by 3 to get 63 ; multiply 63 by 3.77 to get 237.51 ; round 237.51 to 238 ; divide 238 by 63 to get 3.77 $\overline{7}$, which rounds to 3.78 ; and observe that, once again, this mean is inconsistent with the reported sample size. We have made a simple Excel spreadsheet available at https://osf.io/3fcbr that automates the steps of this procedure.

## Practical applications

Using the GRIM technique, it is possible to examine published reports of empirical research to see whether the means have been reported correctly. As psychological journals typically require the reporting of means to 2 decimal places, the sample size corresponding to each mean typically

[^1]needs to be less than 100 in order for its consistency to be checked. However, since the majority of means of interest in experimental psychology are those for subgroups of the overall sample (for example, the numbers in each of two or more experimental conditions), it can still be possible to use the GRIM technique to studies with overall sample sizes substantially above 100 , thus making it applicable to a considerable proportion of published articles ${ }^{4}$.

When an inconsistent mean is uncovered by this method, we of course have no information about the true mean value that was obtained in the study; that can only be determined by a reanalysis of the original data. But such an inconsistency does indicate, at a minimum, that a mistake has been made. When multiple numerical inconsistencies are demonstrated in the same article, we feel the reader is entitled to question what else might not have been reported accurately. And, if the reported test (typically $F$ or $t$ ) statistics and their corresponding $p$ values are themselves perfectly consistent with the reported (inconsistent) means, the authors may well have some explaining to do ${ }^{5}$.

[^2]We now turn to our pilot trial of the GRIM test, in which we examined the consistency of means and other comparable quantities in a substantial sample of recent empirical articles published in leading psychological journals.

## Method

We searched recently published (January 2011 to December 2015; i.e., in the five years preceding the start of our investigations) issues of Psychological Science (PS), the Journal of Experimental Psychology: General (JEP:G), and the Journal of Social and Personality Psychology (JPSP) for articles containing the word "Likert" anywhere in their text. This simple strategy was chosen on the basis that we expected to find Likert-type data reported in most of the articles containing that word. (Some articles that were identified using this method also contained reported means of other, non-Likert type integer data; we also checked the consistency of those means where possible.) We sorted the results in descending order of date of publication (most recent first) and downloaded at most the first 100 matching articles from each journal. Thus, our sample consisted of 100 articles from PS published between January 2011 and December 2015; 60 articles from JEP:G published between January 2011 and December 2015; and 100 articles from $J P S P$ published between October 2012 and December 2015.

We examined the Method section of each study reported in the articles that we had selected to see whether measures suitable for GRIM testing (i.e., those that would generate integer data, or simple composites) were used, and also to determine the sample sizes for the
$p$ values as reported in an article when one or more of the means is known to be inconsistent, this can be a strong indication that the test statistic may be fictitious, having been produced using exactly such a calculation program.
study and, where appropriate, for each condition. A preliminary check was performed by the first author; if he did not see evidence of either GRIM-compatible measures, or any (sub)sample sizes less than 100 , the article was discarded. Subsequently, each author worked independently on the retained articles. We examined the table of descriptives (if present), any other result tables, and the text of the Results section, looking for means or percentages that we could check using the GRIM technique. We noted all of the inconsistent results that we found, doublechecking where necessary that we had correctly identified the per-cell sample size. On the basis of our tests, we assigned each article a subjective "inconsistency level" rating. A rating of 0 (all clear) meant that all the means we were able to check were consistent. We assigned a rating of 1 (minor problems) to articles that contained only one or two inconsistent numbers, where we believed that these were most parsimoniously explained by typographical or transcription errors, and where an incorrect value would have little effect on the main conclusions of the article. Articles that had a small number of inconsistencies, but which might impact the principal reported results of a study, were rated at level 2 (moderate problems); we also gave this rating to articles in which the results seemed to be uninterpretable as described. Finally, we applied a rating of 4 (substantial problems) to articles with a larger number of inconsistencies, especially if these appeared at multiple points in the article. There was no level 3, in keeping with the observation that for a one-item measure such as this, the numerical values are completely arbitrary (Carifio \& Perla, 2007; Murphy \& Reiner, 1984). Finally, individual ratings were compared and differences in ratings resolved by discussion.

## Results

## GRIM analysis of articles

The total number of articles examined from each journal, the number retained for GRIM analysis, and the number to which we assigned each rating, are shown in Table 1. A total of 260 articles were examined. Of these, $189(72.7 \%)$ were discarded, principally because either they reported no GRIM-testable data or their sample sizes were all sufficiently large that no inconsistent means were likely to be detected. Of the remaining 71 articles, 35 (49.3\%) reported all GRIM-testable data consistently and were assigned an inconsistency level rating of 0 (all clear); that is, all of the means (or percentages of integer quantities) that we were able to check were consistent with integers having been divided by the relevant sample sizes (or other denominator), within the limits of rounding.

## Insert Table 1 around here

That left us with 36 articles that appeared to contain one or more inconsistently-reported GRIM-testable data item. Of these, we assigned an inconsistency level rating of 1 (minor problems) to 15 ( $21.1 \%$ of the 71 articles for which we performed a GRIM analysis), 2 (moderate problems) to $5(7.0 \%)$, and 4 (substantial problems) to 16 ( $22.5 \%$ ). In some of these "level 4" articles, over half of the analyzable decimal fractions and percentages were inconsistent with the stated sample size.

## Requests for data

Next, we wrote to the authors of the 21 articles that were rated at level 2 or 4 asking for their data set(s). Our initial e-mail was concise, and stated only that we wished to see these data in order to
allow us to "verify the substantive claims of your article through reanalysis." We took this wording directly (adding only the words "of your article") from article 8.14, "Sharing Research Data for Verification", of the American Psychological Association's (2003) ethics code. In the case of articles published in $J E P: G$ and $J P S P$, we knew that the corresponding author had explicitly agreed to these conditions by signing a copy of a document entitled "Certification of Compliance With APA Ethical Principles" prior to publication.

In response to our 21 initial requests, we received 11 replies. We obtained the requested data sets from eight authors either by immediate return e-mail, or after a request to indicate exactly which studies within an article were relevant; in seven of these cases, the data were complete, while in the eighth, data from one study was not available. One other author immediately promised to send the requested data, but has not done so to date. One author expressed hostility towards the process and sent no further reply to our follow-up e-mails. In one other case, the corresponding author's e-mail address (which was at a popular web-based e-mail hosting service, rather than an institutional address) turned out to have been deleted; in this case, the second author informed us that the first author had left academia and was no longer contactable, and that the location of the data set was unknown.

After two weeks, we sent follow-up requests, including more details about our study and its aims, to the 10 corresponding authors who had not replied to our initial e-mail. In response to these 10 follow-up e-mails, we received eight replies, of which four contained more or less firm offers to send the authors' data sets with varying degree of qualification as to the effort involved (in the end, one of these authors provided us with sufficient information about the data in question to enable us to check the consistency of the means, but we never received anything from the other three), and the other four basically constituted refusals. In these last four cases,
we replied explaining the full details of our method, but we did not receive any further responses. Interestingly, two of the four refusals were identically worded. The remaining two of our 10 follow-up e-mails remained unanswered after ten more weeks had elapsed.

We found the rate of response to our requests for these data sets to be rather disappointing, although we appear to have been more successful than Wicherts, Borsboom, Kats, and Molenaar (2006), who reported that $73 \%$ of the 141 authors they asked had not shared their data after six months. Since the purpose of our request for data sets was to examine the validity of the GRIM technique, rather than to investigate the specific irregularities that might exist in any given article, we did not attempt to emulate the tenacity ${ }^{6}$ shown by Wicherts and colleagues in terms of the number of requests sent per data set, the total time allowed for responses, or the amount of discussion we were prepared to enter into with the authors. We sent a maximum of two unsolicited e-mails to each corresponding author, plus a single follow-up reminder to those authors who offered to share their data but did not, in fact, send these within four weeks.

## Analysis of received data sets

Our examination of the data sets that we did receive showed that the GRIM technique identified one or more genuine inconsistency in each case. We report the results of each analysis briefly here, in the order in which the data sets were received.

[^3]Data set 1. Our GRIM analysis had detected two inconsistent means in a table of descriptives, as well as eight inconsistent standard deviations (the issue of inconsistent SDs will be the subject of a forthcoming article). From the data set, we found that the two inconsistent means and one of the inconsistent SDs were caused by the sample size for that cell not corresponding to the sample size for the column of data in question; five SDs had been incorrectly rounded because the default ( 3 decimal places) setting of SPSS had caused a value of 1.2849 to be rounded to 1.285, which the authors had subsequently rounded manually to 1.29 ; and two further SDs appeared to have been incorrectly transcribed, with values of 0.79 and 0.89 being reported as 0.76 and 0.86 , respectively. All of these errors were minor and had no substantive effect on the published results of the article.

Data set 2. Our reading of the article in this case had detected several inconsistent means, as well as what appeared to be typing mistakes in the reporting of some other statistics and several inconsistently-reported degrees of freedom. The data set revealed that most of these problems were indeed present, and it also showed up a number of other errors in the authors' analysis, such as the use of estimated marginal means reported from an ANOVA, rather than the means and SDs of the original data, to perform post hoc $t$ tests. We subsequently discovered that the article in question had already been the subject of a correction in the journal, which did not address most of the problems that we found. We plan to write to the authors to suggest a number of points that they need to address in a subsequent correction; indeed, a strong case could be made for the entire article to be retracted and resubmitted, as-per the Committee on Publication Ethics guidelines on retraction-the results are probably no longer reliable.

Data set 3. In this case, our GRIM analysis had shown a large number of inconsistent means in two tables of descriptives. The corresponding author provided us with an extensive version of
the data set, including some intermediate analysis steps. We identified that all but one of the columns of data in the two tables of descriptives had been calculated using a formula within Microsoft Excel that included an incorrect selection of cells; for one of the two conditions, this even resulted in the mean and SD of the first condition being included as data points in the calculation of the mean and SD of the second. The corresponding author has assured us that a correction will be issued. It is unclear whether the principal inferential results of the article were affected by these errors; we assume that the authors will verify this in more detail in the course of writing their correction.

Data set 4. In their covering e-mail accompanying their data set, the authors of this article apologized in advance for some possible discrepancies between the sample sizes in the data compared to the article (even though we had not told them that issues with the calculated means formed the reason why we were writing to them). They stated that, due to "a number of computer crashes," they had only been able to retrieve an early version of the data set, and not the final version on which the article was based. We adjusted the sample sizes using the notes that the authors provided, and found that this adequately resolved the inconsistencies in means that we had noticed during our GRIM analysis.

Data set 5. The GRIM analyses in this case found some inconsistent means in the reporting of the data that was used as the input to some $t$ tests, as well as in the descriptives for one of the four conditions in the study. The data set revealed that the former were the result of the authors having reported the $N$ s that were output by SPSS from a $2 \times 2$ repeated-measures ANOVA in which some cases were missing, so that these $N s$ were smaller than the sample sizes that were reported in the method section. The problems in the descriptives were caused by the authors having incorrectly reported the number of participants who met the criteria for exclusion from
analyses for one cell, with the actual number being five larger than the reported value. We were unable to determine to what extent this difference affected the results of the study, although we noted that the per-cell sample sizes were rather small to begin with. We therefore consider that the GRIM analysis made a useful contribution in this case (and the author thanked us warmly for our observations).

Data set 6. In this case, the inconsistencies that we detected with our GRIM analyses turned out to be mostly due to the misreporting by the authors of their sample size. This was not easy to explain as a simple typographical error, as the number of participants was the first word in the first sentence of the methods section, and hence was reported as a word ("Sixty ${ }^{7}$ undergraduates took part"). Additionally, one reported standard deviation that had caused us some concern turned out to have been incorrectly copied and pasted during the drafting process.

Data set 7. This data set confirmed numerous inconsistencies, including several gross errors in the reported degrees of freedom for $F$ tests, from which we had inferred the per-cell sample sizes. Furthermore, a number that was meant to be the result of subtracting one Likert-type item score from another (thus giving an integer result) had the impossible value of 1.5 . We reported this discrepancy and the other inconsistencies to the corresponding author.

Data set 8 . The corresponding author of this study indicated that providing the data set could be problematic, as the data were taken from a much larger longitudinal study. We therefore changed our approach and provided a detailed explanation of the specific inconsistencies we had found, and asked the author to could check these. The author subsequently confirmed that the

[^4]sample size of the study in question had been reported incorrectly, as several participants had been excluded from the analyses but not from the reported count of participants. We were pleased by this result, at least to the extent that it confirmed our calculations, but the author expressed regret at having committed this minor inconsistency and described the exercise as "a good lesson." It did not seem to us that a correction was necessary in this case.

Data set 9. We asked this author for three data sets from a multiple-study article. In one study, we found numerous inconsistencies in a table of descriptives; some of these were explained by missing values for some participants, but others were caused by different numbers than those described in the text having been copied into the table from the SPSS output. In the second, the apparent inconsistencies in the means were caused by missing values for some variables for one participant. The third data set from this article was never obtained.

For completeness, we should also mention that in one of the cases above, the data we received showed that we had failed to completely understand the original article; what we had thought were inconsistencies in the means on a Likert-type measure were due to that measure being a multi-item composite (and correctly reported as such). While our analysis also discovered separate problems with the article in question, this underscores how careful reading is always necessary when using this technique.

## Discussion

We identified a simple method for detecting discrepancies in the reporting of statistics derived from integer-based data, and applied it to a sample of articles published in some of the leading journals in empirical psychology. Of the 71 articles that we were able to test after discarding
those whose sample sizes were too large ${ }^{8}, 36(50.7 \%)$ appeared to contain reporting errors in the summary statistics. Because of the limitations of the GRIM method, we have no way of knowing how many similar inconsistencies might have been discovered in the articles where larger samples of participants were tested, had it been standard practice to report means to three decimal places.

Of the 71 articles that we tested, in 21 cases the inconsistencies seemed to us to be sufficiently serious to warrant asking the authors to share their data. Of these, five essentially refused to share their data, two more did not reply to our requests, and one seems to have become unreachable. Four other authors assented to data sharing but were ultimately uncooperative. Of the nine cases where the authors did share their data (or, in one case, provided technical details about these), we confirmed the existence of reporting problems in all nine, with eight articles showing GRIM-related inconsistencies.

## Some possible objections

[^5]It might be argued that inconsistent means could arise due to truncation being performed when rounding would have been appropriate. For example, with $N=37$ and a total score of 135 , the true mean is $3 . \overline{648}$, which should be rounded to 3.65 , but could conceivably be truncated to 3.64 if the researcher was not paying close attention. While this is perhaps plausible in some cases, we do not consider that it is sufficient to explain most of the occurrences of the problem that can be observed in the published literature, for several reasons. First, most statistical software packages have functionality to either round values to any desired number of decimal places automatically, or to expand the number of decimal places so that borderline cases, such as a third and (currently) last digit of 5, can be corrected. Third, in $50 \%$ of cases (namely, those where the decimal fraction is strictly less than .50 ), truncation gives the same result as rounding. Fourth, as noted earlier, we decided not to consider the (formally incorrect) truncation of exact fractional values ending in 5 at the third decimal place (such as $13 \div 8=1.625$ being truncated to 1.62 ) to be an inconsistency in the present analysis. Fifth, for smaller sample sizes (e.g., those in the fictional case cited at the start of the present article), even an incorrectly truncated value is often still not consistent with the sample size. Corroborating the above analysis, we note that only one of the data sets that we examined contained any cases of inconsistent means due to rounding errors ${ }^{9}$.

Another objection that could be made is that one or more of the reported means could simply have been mistyped at some point during the analysis. However, with modern statistical packages, the means and SDs are typically computed directly by the software from the per-

[^6]participant data as part of the calculation of the test statistic, so that if the combination of the reported means and SDs is consistent with the exact reported $t$ or $F$ value, it is unlikely that the means have been accidentally changed at some stage in the process (as might have happened in the past when the computation of means, SDs, and test statistics typically required the results of several intermediate steps to be noted on paper). Furthermore, in the case of a simple typographical error in copying the value of a mean from the output of the software, one would expect to see a discrepancy between the reported means and SDs on the one hand, and the test statistic on the other. For example, consider the case of two groups of 44 participants, with means (SDs) of 5.36 (1.18) and 4.77 (1.22). The difference between these groups gives $t(86)=2.31, p=.02$. If, when reporting this result, the researcher were to accidentally type 5.63 instead of 5.36 for the first mean, a reader of the subsequent article who applies the GRIM test will detect that this is inconsistent with a sample size of 44 ; but on comparing the means and obtaining $t(86)=3.36, p=.001$, it will be apparent that this inconsistency may well be simply due to a typographical error.

A limitation of the GRIM technique is that, with the standard reporting of means to two decimal places, it cannot reveal inconsistencies with per-cell sample sizes of 100 or more, and its ability to detect such inconsistencies decreases as the sample size increases (or as the number of items in a composite measure increases). However, this still leaves a substantial percentage of the literature that can be tested. Recall that we selected our articles from some of the highestimpact journals in the field; it might well be that lesser journals have a higher proportion of smaller studies. Additionally, it might be the case that smaller studies are more prone to reporting errors (for example, because they are run by laboratories that have fewer resources for professional data management).

A further potential limitation that was raised by one of the corresponding authors with whom we discussed our technique is the case where one or two participants are missing values on individual items in a composite measure. Such missing values could, indeed, introduce false positives into the estimates of inconsistencies. However, in our admittedly modest sample of articles, this issue only caused inconsistencies in one case. We believe that this limitation is unlikely to be a major problem because in any case, as discussed in the previous paragraph, the GRIM test is typically not applicable to measures with a large number of items because of the requirement for the product of the per-cell sample size and the number of items to be less than 100.

## Concluding remarks

On its own, the discovery of one or more inconsistent means in a published article need not be a cause for alarm; indeed, we discovered from our reanalysis of data sets that in many cases where such inconsistencies were present, there was a perfectly innocent explanation, such as a minor error in the reported sample sizes, or some lack of clarity in the reporting of the design of a study. On occasion, too, the GRIM analysis may produce false positives: The reader may have assumed that what looked like a single Likert-type item was in fact a composite measure, or that a stimulus was measured as an integer when in fact it was a continuous real quantity.

It might also be that psychologists are simply sometimes rather careless in retyping numbers from statistical software packages into their articles. However, in such cases, we wonder how many other elementary mistakes have been made in the analysis of the data, and with what effects on the reported results. Indeed, as noted above, in two cases our examination of the data sets we received showed that the authors had made a number of errors in their analyses that were both elementary in their nature and severe in their effects, It is interesting to
compare our experiences with those of Wolins (1962), who asked 37 authors for their data, obtained them from nine authors, and found gross errors in three cases. While the sample sizes here are small, we wonder if some proportion of psychology's replication crisis (Open Science Collaboration, 2015) might be due to the initial (or replication) results being simply the products of erroneous analyses.

Beyond inattention and poorly-designed analyses, however, we cannot exclude that in some cases, a plausible explanation for GRIM inconsistencies is that some form of data manipulation has taken place. For example, in the fictional extract at the start of this article, here is what should have been written in the last sentence:

Participants in the "drunk the Kool-Aid" condition did not report a significantly stronger belief in the ability of pigs to fly $(M=4.79, S D=1.34)$ than those in the control condition $(M=4.27, S D=1.41), t(53)=1.40, p=.17$. In the "published" extract, compared to the above version, the first mean was "adjusted" by adding 0.40 , and the second by subtracting 0.40 . This transformed a non-significant $p$ value into a significant one, thus making the results considerably easier to publish (cf. Kühberger, Fritz, \& Scherndl, 2014).

We are particularly concerned about the eight data sets (out of the 21 we requested) that we believe we will probably never see (five due to outright refusals to share the data, two due to repeated non-response to our requests, and one due to the apparent disappearance of the corresponding author). Refusing to share one's data for reanalysis without giving a clear (typically, ethical) reason is, we feel, professionally disrespectful at best after assenting to such sharing as a condition of publication, as is the case in (for example) APA journals such as JPSP and $J E P: G$. When accompanied by numerical evidence that the results of a published article
may be unreliable, such a refusal will inevitably cause speculation about what those data might reveal. However, throughout the present article, we have refrained from mentioning the titles, authors or any identifying features of the articles in which the GRIM analysis identified apparent inconsistencies. There are three reasons for this; first, the technique as it stands was considered prospective at the time we started to examine the published articles, rather than an established method; second, in any given case, there may be an innocent explanation for any or all of the inconsistencies that we identified in any given article; third, it is not our purpose here to "expose" anything or anyone. We offer our results in the hope that they will stimulate discussion within the field. It would appear, as a minimum, that we have identified an issue worthy of further investigation, and produced a tool that might assist reviewers of future work, as well as those who might wish to check some results in the existing literature.

Acknowledgements
The authors wish to thank Tim Bates and Chris Chambers for their helpful comments on an earlier draft of this article.

References
American Psychological Association (2003). APA ethical principles of psychologists and code of conduct. Retrieved from http://memforms.apa.org/apa/cli/interest/ethics $1 . c f m$

Bakker, M., \& Wicherts, J. M. (2011). The (mis)reporting of statistical results in psychology journals. Behavior Research Methods, 43, 666-678. http://dx.doi.org/10.3758/s13428-011-0089-5

Carifio, J., \& Perla, R. J. (2007). Ten common misunderstandings, misconceptions, persistent myths and urban legends about Likert scales and Likert response formats and their antidotes. Journal of Social Sciences, 3, 106-116. http://dx.doi.org/10.3844/jssp.2007.106.116

Cole, S., Balcetis, E., \& Dunning, D. (2013). Affective signals of threat increase perceived proximity. Psychological Science, 24, 34-40. http://dx.doi.org/10.1177/0956797612446953.

Coombs, C. H. (1960). A theory of data. Psychological Review, 67, 143-159. http://dx.doi.org/10.1037/h0047773

Jamieson, S. (2004). Likert scales: how to (ab)use them. Medical Education, 38, 1212-1218. http://dx.doi.org/10.1111/j.1365-2929.2004.02012.x

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(9): e105825. http://dx.doi.org/10.1371/journal.pone. 0105825

Murphy, K. (Producer), \& Reiner, R. (Director). (1984). This is Spinal Tap [Motion picture]. United States: Embassy Pictures.

Open Science Collaboration (2015). Estimating the reproducibility of psychological science. Science, 349, aac4716. http://dx.doi.org/10.1126/science.aac4716

Reiterman, T., \& Jacobs, J. (1982). Raven: The untold story of the Rev. Jim Jones and his people. New York, NY: E. P. Dutton.

Simonsohn, U. (2013). Just post it: The lesson from two cases of fabricated data detected by statistics alone. Psychological Science, 24, 1875-1888. http://dx.doi.org/10.1177/0956797613480366

Thurstone, L. L. (1927). A law of comparative judgment. Psychological Review, 34, 273-286. http://dx.doi.org/10.1037/h0070288

Wicherts, J. M., Borsboom, D., Kats, J., \& Molenaar, D. (2006). The poor availability of psychological research data for reanalysis. American Psychologist, 61, 726-728. http://dx.doi.org/10.1037/0003-066X.61.7.726

Wolins, L. (1962). Responsibility for raw data. American Psychologist, 17, 657-658. http://dx.doi.org/10.1037/h0038819

| Journal | $P S$ | $J E P: G$ | $J P S P$ | Total |
| :--- | :--- | :--- | :--- | :--- |
| Number of articles | 100 | 60 | 100 | 260 |
| Earliest article date | January 2011 | January 2011 | October 2012 |  |
| Articles with GRIM-testable data | 29 | 15 | 27 | 71 |
| Level 0 articles (no problems detected) | 16 | 8 | 11 | 35 |
| Level 1 articles (minor problems) | 5 | 3 | 7 | 15 |
| Level 2 articles (moderate problems) | 1 | 1 | 3 | 5 |
| Level 4 articles (substantial problems) | 7 | 3 | 6 | 16 |

Table 1 Journals and articles consulted

Notes: $P S=$ Psychological Science. JEP:G = Journal of Experimental Psychology: General.
JPSP: Journal of Social and Personality Psychology.
decimal places.


Note: This figure assumes that means ending in 5 at the third decimal place (e.g.,
Figure 1 Plot of consistent (white dots) and inconsistent (black dots) means, reported to 2
$10 \div 80=0.125$ ) are rounded up; it does not include the possibility, mentioned in the text, of allowing such means to be rounded up or down. This would cause a few extra white dots at sample sizes that are multiples of 8 .

## Appendix A

We show here the e-mails that were sent to the authors of the articles in which we found apparent problems, to request that they share their data. In some cases there were minor variations in wording or punctuation.

The first e-mail, sent in late January 2016:

> Dear Dr. <name>,

We have read with interest your article "<title>", published in <year> in <Journal>.
We are interested in reproducing the results from this article as part of an ongoing project concerning the nature of published data.

Accordingly, we request you to provide us with a copy of the dataset for your article, in order to allow us to verify the substantive claims of your article through reanalysis. We can read files in SPSS, XLS[x], RTF, TXT, and most proprietary file types (e.g., .MAT).

Thank you for your time.
Sincerely,
Nicholas J. L. Brown
PhD candidate, University of Groningen Medical Center
James A. J. Heathers
Postdoctoral Fellow, Poznań University of Medical Sciences

The second e-mail, sent about 10 days after the first if we had received no reply:
Dear Dr. <name>,
Not having received a reply to our first e-mail (see below), we are writing to you again. We apologise if our first message was a little cryptic.

We are working on a technique that we hope will become part of the armoury of peer reviewers when checking empirical papers, which we hope will allow certain kinds of problems with the reporting of statistics to be detected from the text. Specifically, we look at means that do not appear to be consistent with the reported sample size. From a
selection of articles that we have analysed, yours appears to be a case where our technique might be helpful (if we have understood your method section correctly).

However, we are still refining our technique, which is why we are asking 20 or so authors to provide us with data so that we can check that we have fully understood their methods, and see how we should refine the description of our technique to make it as specific and selective as possible. Comparing the results of its application with the numbers in the dataset(s) corresponding to the articles that we have identified will hopefully enable us to understand this process better. So if you could provide us with your data from this article, that would be very helpful.

Kind regards, Nick Brown
James Heathers


[^0]:    ${ }^{2}$ In the studies that we examined for the present article, we sometimes had to make pragmatic assumptions about the reporting of such values. For example, if participants were instructed to make offers of up to five dollars in a dictator game, it is plausible that some would choose an amount such as $\$ 1.50$. On the other hand, when Cole, Balcetis, and Dunning (2013) asked their participants to stand 156 inches ( 13 feet, or approximately 4 meters) from a live tarantula and estimate the distance between them and the spider in inches, it seems unlikely that any responses would have included a fractional component (although, of course, only an inspection of the data set could determine this with certainty).

[^1]:    ${ }^{3}$ An alternative is to adapt the calculation method so that, at the second step, the product of the (unadjusted) sample size and the reported mean is rounded not to the nearest integer, but to the nearest multiple of the granularity of the scale-in this case, 0.33 . However, this is unlikely to be simpler in practice than adjusting the sample size.

[^2]:    ${ }^{4}$ Note that a percentage in the range $0.0 \%-99.9 \%$, even when reported to only one decimal place, can be tested for consistency with a sample size of up to 1000 , as it is, in effect, a fraction reported to three decimal places. For example, if the percentage of a sample of 847 participants endorsing a particular statement is reported as $29.1 \%$, it can be shown that this percentage is inconsistent. The demonstration of this is left as an exercise for the reader.
    ${ }^{5}$ It is quite plausible for a set of correctly reported means and SDs from an article, when input into a simple calculator for $F$ or $t$ tests of the kind that is widely available online, to give a test statistic and $p$ value that are not exactly identical to those reported to two decimal places, because the reported $F / t$ and $p$ values in the article will have been calculated from numbers having greater precision and then rounded. Indeed, if an online calculator gives the exact $F / t$ and

[^3]:    ${ }^{6}$ We note that an investigation into the willingness of researchers to share their data was not the principal goal of these authors either: "Our original aim was to reanalyze these data sets to assess the robustness of the research findings to outliers. We never got that far." (Wicherts et al., 2006, p. 726)

[^4]:    ${ }^{7}$ We have changed the actual number in this example to prevent identification of the article in question. Several other points of detail in the description of the data sets we received have been similarly altered.

[^5]:    ${ }^{8}$ Note that we made the conservative choice to assign an explicit level 0 rating to articles in which all of the means that we were able to test were consistent, even if these means represented only a small percentage of the reported data in the article. Thus, an article reporting six studies, of which five had samples sizes of 500 and one had a sample of 150 divided across three conditions, would receive a level 0 rating if the per-condition means for the one smaller study were consistent. Had we discarded such articles on the basis that the majority of their reported means were untestable, the percentage of articles in which means were tested and found to be inconsistent would have been even higher.

[^6]:    ${ }^{9}$ Specifically, when rounding to three decimal places, SPSS had rounded a value of 8.4448 to
    8.445. The author had then manually rounded this figure to two decimal places as 8.45 , although the correctly reported value should have been 8.44 .

