Durham Research Online

Deposited in DRO:

21 November 2014

Version of attached file:

Published Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Gorard, S. (2015) 'An absolute deviation approach to assessing correlation.', *British journal of education, society and behavioural science.*, 5 (1). pp. 73-81.

Further information on publisher's website:

<http://dx.doi.org/10.9734/BJESBS/2015/11381>

Publisher's copyright statement:

© 2015 Gorard; This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

An Absolute Deviation Approach to Assessing Correlation

Stephen Gorard^{1*}

¹School of Education, Durham University, England.

Author's contribution

The sole author designed, analyzed and interpreted and prepared the manuscript.

Article Information

DOI:10.9734/BJESBS/2015/11381

Editor(s):

- (1) Oyedunni Arulogun, Department of Health Promotion and Education, University of Ibadan, Nigeria.
- (2) Stan Weeber, Faculty of Sociology, McNeese State University in Lake Charles, Louisiana, USA.

Reviewers:

- (1) Perambur S. Neelakanta, CEECS Department, Florida Atlantic University, Boca Raton, Florida 33431, USA.
- (2) Anonymous, National University of Singapore, Singapore.
- (3) Anonymous, Wayne State University, USA.
- (4) Anonymous, George Fox University, USA.
- (5) Anonymous, Università della Calabria, Italy.

Complete Peer review History: <http://www.sciencedomain.org/review-history.php?iid=656&id=21&aid=6222>

Method Article

Received 13th May 2014
Accepted 3rd September 2014
Published 24th September 2014

ABSTRACT

This paper describes two possible alternatives to the more traditional Pearson's R correlation coefficient, both based on using the mean absolute deviation, rather than the standard deviation, as a measure of dispersion. Pearson's R is well-established and has many advantages. However, these newer variants also have several advantages, including greater simplicity and ease of computation, and perhaps greater tolerance of underlying assumptions (such as the need for linearity). The first alternative approach simply divides the co-variance by the mean absolute deviation(s) instead of the standard deviation as in Pearson's R. The second alternative uses the sum of each pair of deviations in x and y instead of the covariance, and again uses the mean absolute deviation(s) as the denominator. All three are compared to one another using 30,000 simulations based on 100 pairs of random numbers. The substantive findings are the same for each approach, and the 'coefficients' correlate with each other (using R) at +0.99 to 1.00. The three approaches also give the same substantive findings when trialled with real-life secondary datasets. This introduction of simpler kinds of correlation forms part of an attempt to simplify the use of numeric analysis, to make it more 'everyday', for the benefit of both analysts and consumers of evidence.

*Corresponding author: Email: s.a.c.gorard@durham.ac.uk;

Keywords: Mean absolute deviation; correlation; quantitative methods initiative; statistical methods; the new statistics.

1. INTRODUCTION

In the UK, there is a long-standing concern about the poor quality and utility of social science and public policy research, and this poor quality has been largely attributed to deficiencies in methods [1,2,3]. Similar concerns have arisen in other countries [4]. Recent capacity-building efforts have focused on a purported lack of work using numbers, and the solution proposed is a national Quantitative Methods Initiative for social science [5]. That initiative is funded by a range of partners including the Higher Education Funding Council for England, the British Academy, the Nuffield Foundation, and the Economic and Social Research Council. The focus is on the development of new methods courses, teaching resources, new training opportunities, work placements, scholarships, and the embedding of 'quantitative' evidence in substantive courses for all upcoming social scientists. The model underlying all of this activity and expenditure is that 'quantitative' work in social science is currently good, especially in economics and psychology, but that it is not always well taught, and there is just not enough of it outside those two areas [6].

Some of this diagnosis may be correct, although it is possible that the problems of social science are as much to with lack of design, genuine curiosity and research integrity as to do with there not being enough studies that involve numbers *per se*. Many of the proposed solutions also may be necessary but they will probably not be sufficient. What is being taught as 'quantitative' methods is currently too often wrong (the meaning of significance tests is widely mis-described, for example – see Gorard 2010), and even more often needlessly complex (the use of multi-level modelling with population data, for example). It is these confusing approaches more than anything that are driving researchers away from the relatively simple use of measurements and frequencies in their work [7,8,9,10]. An alternative approach is to simplify what it means to use numbers in research, to stress the solid links between the logic of analysis for any kind of data, and then try to improve uptake through the pedagogy and persuasion of something like the Quantitative Methods Initiative. This paper is one step in that approach.

This paper introduces two possible alternatives to the more traditional Pearson's R correlation coefficient, both based on using the mean absolute deviation, rather than the standard deviation, as a measure of dispersion. These variants have several advantages, including greater simplicity and ease of computation, and greater tolerance of underlying assumptions (such as the need for linearity). The paper forms part of an attempt to simplify the use of numeric analysis, to make it more 'everyday', for the benefit of both analysts and the consumers of evidence. The approach includes replacing the standard deviation (and its needless squaring and square rooting) with the absolute deviation where possible, introducing a robust absolute deviation effect size [11,12], and absolute deviation regression models, and of course the removal of significance testing and all of its components from everyday analysis [13].

Correlation is a measure of the strength of a relationship between two variables, where each value in one variable has a corresponding or paired value in the other variable. An example might be a list of test scores in maths and reading for the same children. Correlation is the key idea underlying more complex modelling including forms of regression and data reduction. It is usually based on a scatterplot showing a linear or near-linear relationship between the two variables, and the correlation 'coefficient' is an estimate of how well the pairs of values in both variables 'fit' a straight line [14]. In this type of correlation, the values must be real numbers (and there are other techniques for categorical data). The traditional correlation coefficient is Pearson's R.

2. Pearson's R Correlation Coefficient

Correlation generally starts with covariance. 'Covariance' is an estimate of the scale of inter-relatedness of two variables. For a set of cases involving two variables x and y , we can find the deviation of each value from the mean of that variable, multiply the results for each pair of values, and sum these products for all values, before dividing by the number of values. This covariance, or average co-deviation, is large where each value in a pair differs from its mean in the same direction consistently.

$$\frac{\sum_{i=1}^n (x_i - \bar{x}) \cdot (y_i - \bar{y})}{n} \quad (\text{Covariance})$$

The covariance is standardised and adjusted for the scale of measurement by dividing it by the product of the two standard deviations for the two variables. This produces the correlation coefficient known as R.

$$\frac{\frac{\sum_{i=1}^n (x_i - \bar{x}) \cdot (y_i - \bar{y})}{n}}{\sqrt{\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n}} \cdot \sqrt{\frac{\sum_{i=1}^n (y_i - \bar{y})^2}{n}}} \quad (\text{Pearson's R})$$

The numerator is the covariance of the two sets of numbers x (from 1 to n) and y (from 1 to n). This covariance is the sum of: each numbers' deviation from the mean of one set multiplied by the equivalent numbers' deviation from the mean of the other set, all divided by the number of numbers in both sets. The denominator is the product of the standard deviations of the x and y sets of numbers individually. Although usually invisible to analysts using software like SPSS this is quite a complex formula. It is not an easy formula to explain to general under-graduates in social science, for example. The squaring and square-rooting involved in calculating the standard deviation is not intuitive, and has the effect of over-emphasising larger deviations from the mean. There are simpler alternatives.

3. USING THE MEAN ABSOLUTE DEVIATION

Originally the standard deviation (SD) was devised as a way of eliminating the signs for each deviation. By definition, the sum of all of the deviations from the mean for any set of numbers will be zero. Squaring them before summing makes them all positive, and so the sum of all deviations would only be zero if every number in the set was equal to the mean (i.e. there were no deviations).

$$\sqrt{\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n}} \quad (\text{Standard deviation})$$

A simpler alternative to the standard deviation, as a measure of dispersion, is the mean of the absolute deviations from the mean (M|D|). Here the sum is of the absolute values (irrespective of

sign) of each deviation, and because this is divided by the number of values, it is the mean absolute deviation. It gives the same substantive result as the standard deviation because the two are, of course, closely related [12].

$$\frac{\sum_{i=1}^n |x_i - \bar{x}|}{n} \quad (\text{Mean absolute deviation})$$

This is the average by which each value deviates from the mean in each set of numbers (as opposed to the average deviation from the median of the numbers, which could also be used [15]). Such an everyday meaning is lost in the more complex calculation of the standard deviation, because of the squaring and square-rooting. And that is part of the problem when introducing a measure of dispersion like the standard deviation to new researchers. This ease of understanding is only one reason for increasingly preferring the absolute deviation over the standard deviation. The absolute deviation is also more robust when dealing with real-life data not necessarily following a normal distribution [16]. In fact, its use requires no assumptions about the format of the data other than that it is based on real numbers.

There is anyway a widespread but apparently disregarded ambiguity in the use and reporting of standard deviations. The formula for calculating a standard deviation makes it clear that the sum of squares is divided by n, and then the square root of the result is taken. Any square root for a real positive number must always consist of two values (one positive and one negative). In practice, however, the positive square root is often the only one quoted and considered by analysts. Statistical software such as SPSS and office software such as Excel, for example, only ever displays the positive or unsigned result as the standard deviation. So it is only the positive square root that is generally propagated into future calculations (like calculating a standard deviation). The formula should really be adjusted to portray that it is the absolute value of the square root that is being used. But then the main reason why the standard deviation was devised was to evade the use of absolute values, which are inconvenient in many ways for algebraic manipulation. In practice, the use of what are clearly absolute values in the calculation of standard deviations means that any advantage they may have had over absolute deviations now disappears completely.

Using the absolute value of the standard deviation, as happens in practice, can also cause 'errors', since in some calculations half of the possible ensuing outcomes will be ignored. A widespread example is an effect size such as Cohen's *d*, where a difference between means is standardised through division by their standard deviation(s). The result should therefore really be both positive and negative (because the standard deviation in the denominator should be both positive and negative). This would mean that it is not possible to tell from the effect size alone in which direction it operates. For example, it would not portray whether the experimental or control group had performed better in a simple randomised controlled trial. Of course, it is possible to tell this by other means (such as inspecting the data), but the summary effect size itself would be neutral on the matter, even using the convention of putting an intervention group before the control group score. Therefore, a simpler form of analysis based solely on absolute deviations could be presented instead. This is the idea examined further in this paper.

4. AN ABSOLUTE DEVIATION CORRELATION COEFFICIENT

One simpler alternative for a correlation coefficient would be to replace the standard deviations in the denominator of *R* with the mean absolute deviation of each set of numbers instead. This would be:

$$\frac{\sum_{i=1}^n (x_i - \bar{x}) \cdot (y_i - \bar{y}) / n}{\sum_{i=1}^n |x_i - \bar{x}| / n \cdot \sum_{i=1}^n |y_i - \bar{y}| / n} \quad (\text{new correlation coefficient RA1})$$

This approach does not involve squaring the deviations of each value from the mean, summing the results and then finding the square root. The resulting formula is therefore simpler than for Pearson's *R*. It is the sum of the products of each pair of deviations divided by the product of the sums of each pair of deviations. As a measure of dispersion, the mean absolute deviation is robust and efficient. Unlike the standard deviation, it does not artificially and unnecessarily inflate the larger deviations by squaring them, and then incompletely square-rooting the results. In almost all other practical circumstances it does not matter whether the standard or absolute deviation is used. For

example, creating 30,000 pairs of samples, each of 100 random numbers, allows the correlation between each pair to be assessed. Doing this with both Pearson's *R* and the absolute deviation correlation (here termed 'RA1') yields two sets of 30,000 examples Fig. 1. This shows that there is a clear and straight-line equivalence between the two measures of correlation. Using *R*, their inter-'correlation' is reported as +1.0.

A standard deviation is generally larger than an absolute deviation for the same set of figures (because it inflates the figures slightly via the squaring and square-rooting process). This means, in turn, that the absolute deviation correlation coefficient (created by dividing the covariance by the absolute deviations) is generally larger in absolute terms than *R* for the same two sets of figures. RA1 is usually greater than *R* by a factor of around 1.3, because the absolute deviation is usually smaller than the standard deviation by a factor of about 0.77. Some commentators have stated that there is a simple conversion from SD to *M|D|*, through multiplying by $\sqrt{2/\pi}$ or 0.77. In fact this is not so. Each SD can have more than one *M|D|*, and vice versa, again caused by the process of squaring the deviations, summing and then square-rooting. And this is why the graph in Fig. 1 appears to diverge somewhat from a straight line towards the extremes. Nevertheless, either coefficient will work as well as the other in practice, and the 0.77 or 1/0.77 conversion will usually work well enough to see what the other coefficient would have been. Before comparing the strengths and weaknesses of each approach more fully, it is interesting to see if the coefficient can be simplified even further.

5. AN ADDITIVE COEFFICIENT

In some respects it seems obvious to create a correlation coefficient that matches Pearson's *R* but with the standard deviations replaced by absolute deviations (as above). It might also be possible to create a kind of correlation coefficient in which the deviations for each pair of numbers are summed, and are divided by the sum of the sums of the deviation for each number separately. Where both numbers in a pair deviate in the same direction they will add together, and when in the opposite direction they will tend to cancel out. The sum of these combined deviations for both pairs together can be assessed as a proportion of the sum of the absolute deviations for each set of numbers alone.

Fig. 1. Scatterplot of values for R (x axis) and RA1 (y axis), each based on 30,000 trials using two sets of 100 random numbers

$$\frac{\sum_{i=1}^n |(x_i - \bar{x}) + (y_i - \bar{y})|/n}{\sum_{i=1}^n |x_i - \bar{x}|/n + \sum_{i=1}^n |y_i - \bar{y}|/n} \text{ (new correlation coefficient RA2)}$$

And, unlike the more complex Pearson's R, the value of n can be cancelled in numerator and denominator. This will yield:

$$\frac{\sum_{i=1}^n |(x_i - \bar{x}) + (y_i - \bar{y})|}{\sum_{i=1}^n |x_i - \bar{x}| + \sum_{i=1}^n |y_i - \bar{y}|} \text{ (new correlation coefficient RA2)}$$

This is even simpler again, and therefore should be even easier for new researchers to learn and use. As with RA1 the results from using RA2 again correlate very highly with those for Pearson's R. Here the results are again illustrated for 30,000 pairs of samples, each of 100 random numbers. Running this simulation with both Pearson's R and the absolute deviation correlation coefficient (RA2) yields two sets of 30,000 examples Fig. 2. This shows that there is a clear and straight-line equivalence between the two measures of correlation. Using R, their 'correlation' is reported as over +0.99.

The use of absolute values in both top and bottom of the formula means that RA2 is always positive (or zero). Zero for RA2 is equivalent to -1 for R. Zero for R is equivalent to 0.667 for RA2 (where the regression line crosses the y axis in Fig. 2. When R is 1, RA2 is also 1. What this might mean is discussed in the final section.

6. A REAL-LIFE COMPARISON

It is also interesting to observe the behaviour of these two new proposed correlation approaches using real datasets. One such analysis is based on the annual school-level census of secondary schools in England. The figures for the characteristics of pupils in each school can be used to assess how clustered within specific schools, or 'segregated' between schools, are those pupils with any indicator of potential disadvantage. This is done using a segregation index, and the detailed method is reported elsewhere [17]. This means that in each year it is possible to say how segregated the school system is in terms of poor children (those eligible for free school meals or FSM), children with a disability or special educational need (SEN), and children from a minority ethnic background (non-White). It is clear that the levels of segregation and the changes over time are different for these different indicators [18]. While segregation by poverty has risen since 1998 and is now falling again, segregation of SEN pupils has dropped over time and is now rising slightly, and

segregation by ethnic origin has fallen continuously. This picture yields the correlations in Table 1.

Table 1. Comparing three correlation 'coefficients' for the same datasets: correlating the unintended segregation of FMS, SEN and non-white pupils in secondary schools in England, 1998-2014

Correlation approach	R	RA1	RA2
Segregation FSM and SEN pupils	-0.85	-1.09	0.26
Segregation SEN and non-white pupils	+0.61	+0.79	0.94

Table 1 shows that each coefficient remains related to each other in the same way as in the simulations above. R and RA1 are positive or negative together, and R is between 0.77 and 0.78 times RA1 for both figures (see above). They are both measuring the same thing and yielding the same substantive results. RA2 does not have actual negative values, but as Fig. 2 shows any value below 0.66 can be regarded as showing a negative relationship. Both values of RA2 in Table 1 could be read from the corresponding value of R in Fig. 2. Again RA2 is yielding the same substantive results as R, but perhaps is a slightly harder to interpret form.

7. COMPARING THE THREE APPROACHES

What this brief paper has shown is that it is possible to create alternatives to Pearson's R that are to some extent simpler (algebraically), based on the absolute mean deviation [19]. In practice, and as long as all three coefficients are understood, which one is used in practice would make little difference (and many more trials have been conducted than are reported in the Figures and Tables here). They all give the same, or very nearly the same, substantive result in the contexts presented here. In many respects also they all have similar characteristics.

Pearson's R has the advantages of being built into software, its current use in other contexts, and its familiarity for existing users, especially concerning its limits, scale and importance in any context. The latter is important given that significance testing and its derivatives are inappropriate in most real-life research contexts [20], and do not provide an accurate probability of the 'significance' of a correlation even when all assumptions are met [21]. What is needed instead is judgement based on the scale of the

correlation (the effect size), the number of cases, and the quality of data [22]. Familiarity helps such judgement. R also has the advantage over RA2 of having a zero that portrays easily when there is no linear covariance. It is symmetrical, as long as it is remembered that this is based on using the absolute values of the standard deviations instead of the two (both positive and negative) versions of the standard deviations themselves. Perhaps most importantly, it is clearly and conveniently bounded by -1 and +1.

RA1 is very similar to R as would be expected where the covariance is being divided by two closely related estimates of dispersion. Clearly, when x and y are unrelated the covariance or denominator for both coefficients is zero, and thus both coefficients are zero. Therefore RA1 shares this advantage with R. However, RA1 is generally simpler to explain to new researchers, and slightly easier to calculate. Based on the mean absolute deviation, it will be more robust in the majority of real-life situations where the data contains errors and/or is not perfectly normally-distributed [16]. Where the dataset consists of ordinal values rather than real numbers, Spearman's rank correlation coefficient is a useful alternative having the distribution-free appeal of RA1 and the convenience of R [23]. However, using Spearman's rho where the dataset consists of real numbers means discarding a considerable amount of information – treating real numbers as mere ranks – which also means that the results for rho correlate less highly with R than those for RA1 and RA2 do.

RA1 is less distorted than R by larger deviations due to the lack of squaring, and so does not promote the deletion of inconvenient data by unwary analysts. It also appears more tolerant of non-linearity. As a scatterplot diverges from a linear crossplot of x and y but the two sets of data retain a clear relationship, such as when one is a power of the other, R decreases rapidly. In these circumstances, RA1 does not decrease. The real-life correlation of x with x+1 is no stronger than the correlation of x with x², for example. R merely makes it appear to be so, because x is not clearly *linearly* related with x². This stable and robust characteristic of RA1 is one reason why the factor of 1.3 (above) is only a guide for when x and y are linearly related. Using RA1 early in an investigation might help pick up indications of a more complex relationship than a simple linear one between x and y.

Fig. 2. Scatterplot of values for R (x axis) and RA2 (y axis), each based on 30,000 trials using two sets of 100 random numbers

RA2 is different to both R and RA1 in some respects. Based on addition not multiplication, and on absolute values in both numerator and denominator, it does not have a negative and positive form. In this respect it is no different to traditional effect sizes calculated properly (see above). It might therefore be necessary to make a cursory examination of the data or look at a crossplot in order to be certain of the direction of any relationship. The major unique advantage of RA2 is its conceptual and arithmetic simplicity. Like RA1 it is more robust, less distorted by large deviations, and more tolerant of non-linearity than R.

Where the results from each approach differ somewhat, as they do at extremes in Fig. 1 and throughout Fig. 2, it is not immediately clear which result is preferable. To the extent that the different coefficients yield the same substantive results (as in measuring a temperature in either Centigrade or Fahrenheit) then scientifically it does not matter which one is used. Where they differ, given that the distortion caused by squaring is already well-known, there is an argument that the absolute deviation coefficients might more accurately portray what is being estimated.

8. CONCLUSION

Clearly more work would be needed than presented in this introductory paper to make either RA1 or RA2 something that could be used safely in practice, and offered as a genuinely simpler alternative to R. Analysts would also need more practice in assessing the meaning of a new coefficient value. It is not as easy, for example, as using the mean absolute deviation to create a new effect size [11]. But the work so far shows that in principle the mean absolute deviation can be used instead of the standard deviation in a variety of settings, probably throughout descriptive statistics [12]. The overall objective is to make working with numbers as simple as possible for reluctant social scientists. At present, in the UK at least, considerable resources are being spent on trying to widen the use of 'quantitative' methods. Almost all of the effort lies in persuasion and pedagogy. These might be necessary, but will not be sufficient in themselves. This paper is part of an attempt to point in a different way – correcting the logical errors and simplifying what it is deemed necessary to know, before then improving the teaching of the corrected and simpler methods.

COMPETING INTERESTS

Author has declared that no competing interests exist.

REFERENCES

- McIntyre D, McIntyre A. Capacity for research into teaching and learning, Swindon: Report to the ESRC Teaching and Learning Research Programme; 2000.
- Schuller T. Capacity building in educational research: Sketching an international picture; 2007. accessed 23/5/14. Available: <http://www.scotedreview.org.uk/pdf/198.pdf>.
- Goldacre B. Building evidence into education, DfE, 2013, Accessed 23/5/14, Available: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/193913/Building_evidence_into_education.pdf
- National Science Foundation Scientific research in education. Washington DC: National Academy Press; 2002.
- MacInnes J. Proposals to support and improve the teaching of quantitative research methods at undergraduate level in the UK. Report of the Strategic Adviser for Quantitative Methods to the Economic and Social Research Council; 2009. Accessed 180714. Available: http://www.esrc.ac.uk/images/undergraduate_quantitative_research_methods_tcm8-2722.pdf
- ESRC (2014) Quantitative Methods Initiative. Accessed 21/06/14. Available: <http://www.esrc.ac.uk/research/skills-training-development/qmi/index.aspx>.
- Cooper L, Shore F. Students' misconceptions in interpreting center and variability of data represented via histograms and stem-and-leaf plots. Journal of Statistics Education. 2008;16(2):13.
- Hampel F. Is statistics too difficult? Research Report 81, Seminar fur Statistik, Eidgenossische Technische Hochschule, Switzerland; 1997.
- Murtonen M, Lehtinen E. Difficulties experienced by education and sociology students in quantitative methods courses. Studies in Higher Education. 2003;28(2):171-185.
- Watts D. Why is introductory statistics difficult to learn? The American Statistician. 1991;45(4):290-291.
- Gorard S. The possible advantages of the mean absolute deviation 'effect' size, Social Research Update. 2013;65(Winter 2013):1-4. Available: <http://sru.soc.surrey.ac.uk/SRU64.pdf>.
- Gorard S. The widespread abuse of statistics by researchers: What is the problem and what is the ethical way forward? Psychology of Education Review. 2014;38(1):3-10.
- Gorard S. Introducing the mean absolute deviation 'effect' size. International Journal Research and Methods in Education. 2014;107:4.
- Gorard S. Quantitative methods in social science: The role of numbers made easy, London: Continuum; 2003.
- El Amir E. Mean absolute deviation about median as a tool of explanatory data analysis, International Journal of Research in Arts and Social Sciences. 2012;11:3. Available: http://www.arpapress.com/Volumes/Vol11Issue3/IJRRAS_11_3_17.pdf.
- Gorard S. Revisiting a 90-year-old debate: The advantages of the mean deviation. The British Journal of Educational Studies. 2005;53(4):417-430.
- Gorard S, See BH. Overcoming disadvantage in education, London: Routledge; 2013.
- Gorard S, Hordosy R, See BH. Narrowing the determinants of segregation between schools 1996-2011. Journal of School Choice. 2013;7(2):182-195.
- El Amir E. On uses of mean absolute deviation: Decomposition, skewness and correlation coefficients, Metron: International Journal of Statistics. 2012;LXX:2-3.
- Berk R, Freedman D. Statistical assumptions as empirical commitments; 2001. Accessed 030714. Available: <http://www.stat.berkeley.edu/~census/berk2.pdf>.
- Gorard S. All evidence is equal: The flaw in statistical reasoning. Oxford Review of Education. 2010;36(1):63-77.

22. Lipsey M, Puzio K, Yun C, Hebert M, Steinka-Fry K, Cole M, Roberts M, Anthony K, Busick M. Translating the statistical representation of the effects of education interventions into more readily interpretable forms. Washington DC: Institute of Education Sciences; 2012.
23. Siegel S. Nonparametric statistics for the behavioural sciences, Tokyo: McGraw Hill; 1956.

© 2015 Gorard; This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Peer-review history:
The peer review history for this paper can be accessed here:
<http://www.sciencedomain.org/review-history.php?iid=656&id=21&aid=6222>