



ELSEVIER

Contents lists available at ScienceDirect

Journal of Informetrics

journal homepage: www.elsevier.com/locate/joi

Regression for citation data: An evaluation of different methods



Mike Thelwall*, Paul Wilson

Statistical Cybermetrics Research Group, School of Mathematics and Computer Science, University of Wolverhampton, Wulfruna Street, Wolverhampton WV1 1LY, UK

ARTICLE INFO

Article history:

Received 2 July 2014
 Received in revised form
 26 September 2014
 Accepted 30 September 2014
 Available online 22 October 2014

Keywords:

Informetrics
 Altmetrics
 Citation distributions
 Lognormal
 Powerlaw
 Regression

ABSTRACT

Citations are increasingly used for research evaluations. It is therefore important to identify factors affecting citation scores that are unrelated to scholarly quality or usefulness so that these can be taken into account. Regression is the most powerful statistical technique to identify these factors and hence it is important to identify the best regression strategy for citation data. Citation counts tend to follow a discrete lognormal distribution and, in the absence of alternatives, have been investigated with negative binomial regression. Using simulated discrete lognormal data (continuous lognormal data rounded to the nearest integer) this article shows that a better strategy is to add one to the citations, take their log and then use the general linear (ordinary least squares) model for regression (e.g., multiple linear regression, ANOVA), or to use the generalised linear model without the log. Reasonable results can also be obtained if all the zero citations are discarded, the log is taken of the remaining citation counts and then the general linear model is used, or if the generalised linear model is used with the continuous lognormal distribution. Similar approaches are recommended for altmetric data, if it proves to be lognormally distributed.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

The use of performance monitoring for university research has increased over the past few decades. This is most evident in national research evaluation exercises, such of those in the UK (Mryglod, Kenna, Holovatch, & Berche, 2013), Australia (ARC, 2014), New Zealand (Anderson, Smart, & Tressler, 2013) and Italy (Abramo, D'Angelo, & Di Costa, 2011). This climate not only affects the allocation of research funding in many cases but can also change the behaviour of individual researchers as they come to terms with the assessment system (Butler, 2003). Although the most important performance monitoring exercises often rely on peer review, both the UK (REF, 2013) and Australia (ARC, 2014) consider citations for some subject areas, and there are advocates of increasing use of citations for some types of science when the results correlate because citation metrics are much cheaper than peer review (Abramo, Cicero, & D'Angelo, 2013; Franceschet & Costantini, 2011; Mryglod et al., 2013), although no simple method is likely to work (Butler & Mcallister, 2011). In addition, citations are used for formal and informal evaluations of academics (Cole, 2000) and the Journal Impact Factor (JIF) is a widely recognised and used citation metric.

The range of metrics of relevance to science has recently increased with the emergence of webometrics (Almind & Ingwersen, 1997), which includes a range of new impact indicators derived from the web (Kousha & Thelwall, 2014) and

* Corresponding author. Tel.: +44 1902 321470.

E-mail addresses: m.thelwall@wlv.ac.uk (M. Thelwall), PaulJWilson@wlv.ac.uk (P. Wilson).

altmetrics (Priem, Taraborelli, Groth, & Neylon, 2010), which incorporate many attention and impact indicators derived from social web sites (Priem, 2014). Altmetrics seem particularly promising to help researchers to identify recently published articles that have attracted a lot of attention (Adie & Roe, 2013) and to give evidence of non-standard impacts of research that can be added to CVs (ACUMEN, 2014; Piwowar & Priem, 2013). Statistical analyses of some of these have started to generate new insights into how science works (Mohammadi & Thelwall, 2014; Thelwall & Maflahi, *in press*) and the types of research impacts that are not recognised by traditional citation counts (Mohammadi & Thelwall, 2013).

Because of the many uses of citations within science, it is important to understand as much as possible about why they are created and why one article, researcher or research group may be more cited than another. Whilst citations may be given to acknowledge relevant previous work (Merton, 1973), they can also be used to criticise or argue against it (MacRoberts & MacRoberts, 1996) and so citations are not universally positive. Moreover, citations do not appear to be chosen in a dispassionate, unbiased way (Borgman & Furner, 2002). For example, researchers in some fields tend to cite papers written in the same language (Yitzhaki, 1998), highly relevant citations may be systematically ignored (McCain, 2012) and fame seems also to attract citations (Merton, 1968). There are also field differences in the typical number of citations received by papers (e.g., Waltman, van Eck, van Leeuwen, Visser, & van Raan, 2011) and review articles are more likely to be highly cited than other articles (e.g., Aksnes, 2003). From a different perspective, factors that associate with highly cited papers, such as collaboration, internationality and referencing patterns (Didegah & Thelwall, 2013b; Sooryamoorthy, 2009), are important because they may push researchers and funders towards more successful types of research.

Although some of the factors affecting citations discussed above have been discovered using qualitative methods, such as interviews with authors, statistical methods are needed to identify the magnitude of factors and to detect whether they apply in particular cases. The simplest approach is probably to compare the average number of citations (or citation-based indicators) for one set of papers against that of another to see which is higher (van Raan, 1998). Another strategy is to assess whether citation counts correlate significantly against another metric that is hypothesised to be related (Shema, Bar-Ilan, & Thelwall, 2014). The most discriminating methods used so far are regression-based because they allow the effects of multiple variables to be examined simultaneously. In particular, regression guards against one factor being identified as significant (e.g., international collaboration) when another factor (e.g., collaboration with the USA) is underlying cause of higher (or lower) citations.

There is no consensus about the best regression method for citation data. Methods used so far include ordinary least squares linear regression (Aksnes, Rørstad, Piro, & Sivertsen, 2013; Dragos & Dragos, 2014 [citations per publication used as the dependant variable]; Foo & Tan, 2014; He, 2009; Mavros et al., 2013; Rigby, 2013 [adding 1 to citations, dividing by a time normalised value and taking their log]; Tang, 2013 [adding 1 to citations and taking their log]; Stewart, 1983), logistic regression (Baldi, 1998; Bornmann & Williams, 2013; Kutlar, Kabasakal, & Ekici, 2013; Sin, 2011; Willis, Bahler, Neuberger, & Dahm, 2011; Xia & Nakanishi, 2012; Yu, Yu, & Wang, 2014), a distribution-free regression method (Peters & van Raan, 1994), multinomial logistic regression (Baumgartner & Leydesdorff, 2014) and negative binomial regression (Chen, 2012; Didegah & Thelwall, 2013a, 2013b; McDonald, 2007; Thelwall and Maflahi, *in press* [for altmetrics]; Walters, 2006; Yoshikane, 2013 [for patent citations]).

The typical distribution of citations is highly skewed (de Solla Price, 1965; Seglen, 1992), so that tests based upon the normal distribution (e.g., ordinary least squares regression) are not appropriate if the data is raw citation counts. Logistic regression can avoid this issue by predicting membership of the highly cited group of papers rather than directly predicting citations. Whilst negative binomial regression can cope with skewed data and is designed for discrete numbers (Hilbe, 2011), the most appropriate distribution for citations to a collection of articles from a single subject and year seems to be the discrete lognormal distribution (Evans, Hopkins, & Kaube, 2012; Radicchi, Fortunato, & Castellano, 2008; Thelwall & Wilson, 2014) and the hooked power law is also a reasonable choice (Thelwall & Wilson, 2014). Although many articles suggest a power law for the tail of citation distributions (e.g., Yao, Peng, Zhang, & Xu, 2014) this is not helpful for statistical analyses that need to include all cited articles and is broadly consistent with a lognormal distribution for all articles, although small discrepancies may be revealed by fine-grained analyses (Golosovsky & Solomon, 2012). Citations to articles from individual journals almost always conform to a lognormal distribution (Stringer, Sales-Pardo, & Amaral, 2010), as do some other citation-based indicators also follow a lognormal distribution (e.g., generalised *h*-indices: Wu, 2013). Although it has not been fully tested, it seems likely that most sets of articles from a specific time window will approximately follow a discrete lognormal distribution, unless the time window is too long or very recent. Hence it is not clear that negative binomial regression is optimal when the dependant variable is citation counts.

Neither the discrete lognormal or the hooked power law distributions have been used for regression because it seems that no software exists for this. An alternative strategy would be to take the log of the citations and then use the general linear model to analyse them with the assumption that the logged citations would be normally distributed (the general linear model assumes that the error terms or residuals are normally distributed). Although the log of the continuous version of the lognormal distribution is a perfect normal distribution, the same is not true for the discrete lognormal distribution and so it is not obvious that this will work. Moreover, the use of log transformation for citation data has been argued against for classification purposes because of the variance reduction that it introduces (Leydesdorff & Bensman, 2006), but this is not evidence that it will not work for regression. This article assesses both of these and the continuous normal distribution in order to identify the most powerful, regression-based approach for citations and similar data, such as altmetrics. The results will help to ensure that future statistical analyses of the factors affecting citation counts are as powerful and reliable as possible.

2. Citation distributions and statistical tests

Statistical regression models attempt to identify the relationship between a dependant variable (citations in this case) and one or more independent variables (e.g., the field in which an article is published or the number of authors). The relationship is typically identified by an algorithm that estimates parameters for the independent variables in order to generate a model that, for any given set of their values, predicts an expected value for the dependant variable. For example, if citations were regressed against scientific field, then, for any given field, the equation would predict the expected number of citations for that field. Moreover, tests associated with the regression could confirm or reject the hypothesis that two or more fields have different means. The strategy for fitting the model depends upon the distribution of the error terms, which are the difference between the predicted values and the observed values for the data. If these error terms are normally distributed then the regression can use the general linear (ordinary least squares) model (Garson, 2012). If the error terms are not normally distributed, however, then using the general linear model can give misleading results. To illustrate this, since citations are highly skewed, if the general linear model were to be used to fit an equation then the estimates made by the equation would be greatly influenced by the few highly cited articles because the normal distribution assumption would lead to a prediction of a higher estimated mean to reduce the otherwise extremely low chance of the high citation. In other words a citation that is high enough to be an outlier for the normal distribution could still be within the expected range of values for the lognormal distribution.

Switching to the generalised linear model (Dobson & Barnett, 2008) can resolve the problem of error terms not being normally distributed when fitting the general linear model. This allows the distribution of error terms to be specified as something other than normal, such as negative binomial, Poisson or lognormal, and fits the regression equation using maximum likelihood estimation (Millar, 2011). For example, if the distribution is specified as negative binomial, which is skewed, then the regression equation fitted to a citation distribution will be much less influenced by high citation scores than in the case of the general linear model, because high values are more likely for the negative binomial distribution. Hence negative binomial regression (Hilbe, 2011) seems to be a better choice than does general ordinary least squares linear regression.

The negative binomial model assumes that the error terms (i.e., residuals) follow a negative binomial distribution. Although this is a skewed distribution, it is not the same as the lognormal distribution and hence is unlikely to fit lognormal data well, although it is not clear whether the difference between the two approaches would be significant in practice. Here, a discrete version of the lognormal distribution would ideally be needed to model the error terms for the generalised linear model. Unfortunately, however, whilst many different versions of the generalised linear model are available in statistical software, such as R, SPSS and STATA, at the time of writing (June 2014) none includes either the discrete lognormal distribution and so this is not a practical option. Note that although R includes the Poisson lognormal distribution (Aitchison & Ho, 1989), this is different from the discrete lognormal distribution. The GAMLSS package in R can be used to fit models based upon the continuous lognormal distribution (Stasinopoulos & Rigby, 2007) with the generalised linear model and so this is another logical alternative.

An alternative to generalised linear regression with a lognormal error function is to take the log of the discrete lognormal data and analyse it with the general linear model with the assumption that the error distribution, despite being derived from a discrete version of the lognormal distribution, will be close enough to normal for models to be effectively fitted. This is plausible because large values have the most influence on model fitting and the log function eliminates these. Nevertheless, there is a problem in that the log function is undefined for zero citations. To circumvent this, either all zero citations should be removed, or an offset of 1 should be added to all citation counts (a common statistical method, e.g., Tabachnick & Fidell, 2001, p. 81) to allow the logs to be calculated (Tang, 2013). Neither of these is an ideal solution since both introduce a systematic source of bias into the data. The first method (removing all zero citations) seems to be the most problematic, however, since it involves truncating one side of the distribution, which will bias the regression towards estimating higher citation means and may affect factors associating with lower citation scores more than factors associating with higher citation scores, reducing the power of the test. In addition, the truncation reduces the sample size and hence also the power of the test.

Finally, it is possible, although it seems unlikely, that citations do not follow a lognormal distribution but instead the citations to a collection of articles follow another distribution but with factors that combine the other distribution in a way that gives a combined distribution that is lognormal. Whilst theoretically possible, this is unlikely given the range of different sets of articles that have been tested and found to be approximately discrete lognormal (Thelwall & Wilson, 2014).

3. Research question

In the absence of software that uses a discrete version of the lognormal distribution in the generalised linear model, this article assesses three logical alternatives, driven by the following research question: Which of the following is the most appropriate for discrete lognormal data in the sense that it does not give an inflated chance of false positive results whilst being powerful enough to distinguish minor factors in the data?

1. Negative binomial regression for the raw data.
2. Regression with the generalised linear model with a continuous lognormal distribution for the error terms after removing zeroes.

3. Regression with the general linear model on the log of the raw data after removing zeroes.
4. Regression with the generalised linear model with a continuous lognormal distribution for the error terms after adding 1 to the observations.
5. Regression with the general linear model on the log of the raw data after adding 1 to the observations.

4. Methods

In order to address the research question, the three regression approaches must be applied to citation-like data with known relationships. This can be achieved by tests on simulated discrete lognormal data with and without relationships. Although there are infinitely many combinations of dependant variables that could be tested, the simplest case will give the clearest results and so only one relationship will be tested: that of a binary factor so that the two values of the factor correspond to different means for the lognormal distribution (but the same standard deviation).

Datasets were randomly created with citations, C , and a binary factor, B , taking two values, 0 and 1, so that the population distribution was approximately lognormal (due to discretization) with mean μ_{\equiv} and standard deviation σ . For simplicity, half of each sample was taken for $B=0$ and half for $B=1$. A range of differences between μ_0 and μ_1 was tested to see how large the difference must be in order to reliably detect the binary factor. This is related to the power of the test. The following R commands were used for all of tests reported here.

```
glm.nb(C ~ f, data = all_data) #negative binomial
gamlss(C ~ f, family = LOGNO(), data = zero.truncated_data) #continuous lognormal
aov(Clog ~ f, data = zero.truncated_data) #ANOVA on log of C
gamlss(C.plus1 ~ f, family = LOGNO(), data = all_data) #continuous lognormal +1
aov(C.plus1.log ~ f, data = all_data) #ANOVA on log of C +1
```

To test for false identification of factors, data was drawn as above except that $\mu_0 = \mu_1$. If, despite the identical means, a factor is detected by regression then this is a false positive result.

The test procedure was to generate 1000 random datasets for each set of parameters to be tested and to calculate the proportion of the datasets giving correct results for each of the above five models. This procedure was repeated for a variety of different sample sizes, means and standard deviations to reduce the chance of the results being specific to the set of parameters chosen.

5. Results

The five approaches were first fitted to discrete lognormal data without any factors using 1000 simulations at each of a range of different sample sizes. As shown in Fig. 1, when the log of the mean is 0.5 and the log of the standard deviation is 1, both ANOVA variants and both continuous lognormal models incorrectly identify factors at the approximately the rate corresponding to the level of significance used (0.05), which is the desired behaviour and suggests that the deviation from normality of the logged discrete data does not cause a problem for the tests. In contrast, and as suggested by the discussion above, negative binomial regression incorrectly identifies a factor at about triple the rate of that suggested by the confidence level. When the log of the standard deviation is increased to 2, the ANOVA and continuous lognormal models keep the same approximately 5% failure rate but negative binomial regression identifies non-existent factors about half of the time and is less successful as the sample size increases (Fig. 2). The same patterns also occur for log means of 1 and 1.5 and so the results do not seem to be dependent on the selected mean.¹

The regressions against simulated discrete lognormal data with a simple factor give information about the power of the tests to correctly identify factors. When the log of the means of the two values of the factor differ by only 0.05 (i.e., one log mean is 0.50 and the other is 0.55) all of the methods need larger sample sizes to increase their chance of detecting it (Fig. 3). The most powerful of the three methods is negative binomial regression and the least powerful are the standard ANOVA and the standard continuous lognormal model. The apparent additional power of the negative binomial over that of the standard ANOVA and continuous lognormal is about the same as the difference in Fig. 1, suggesting that the additional positives are due to the tendency to generate false positive results rather than it being a more powerful method. The additional power of the ANOVA + 1 and lognormal + 1 tests are presumably due to the increased sample size and the increased difference between the two factors caused by not deleting the zeroes. Although Fig. 3 is based on a single mean and standard deviation, the pattern is almost identical for log means of 1 or 1.5, and the relationship between the methods is the same if the log standard deviation is doubled to 2 (Fig. 4), although this makes the factor harder to detect.

If the factor is easier to detect because there is a higher difference in log means (e.g., 0.1 rather than 0.05) then (Fig. 5) then the advantage of the modified ANOVA and lognormal over the standard ANOVA and continuous lognormal is more pronounced.

¹ The full code for the tests, together with all the data used and additional graphs (including any marked as “not shown” in the text) is available at http://figshare.com/articles/Data_from_the_paper_Regression_for_Citation_Data_/1090729.

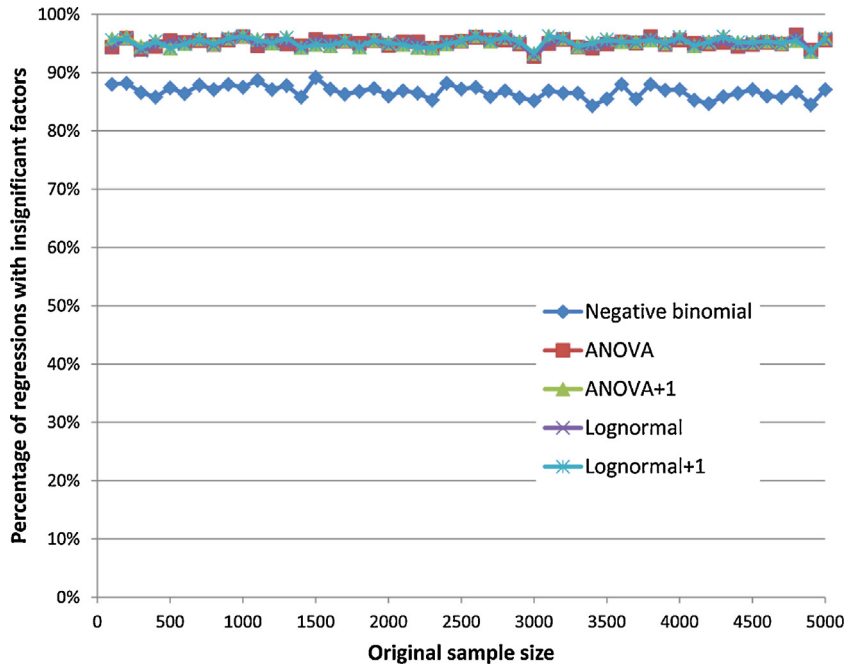


Fig. 1. Regressions against simulated discrete lognormal data (log mean: 0.5, log standard deviation: 1) at different sample sizes without factors ($n = 1000$ iterations). All lines overlap except for the negative binomial line.

6. Discussion and conclusions

A limitation of the tests reported here is that they only deal with the simplest case of a single factor but it seems likely that the same conclusions would be drawn for more complex factors. Moreover, the unreliability of negative binomial regression for citations is probably overestimated by the simulation approach used here since real citation data will be bounded, especially if the data is from a small set of years that are not a long time in the past.

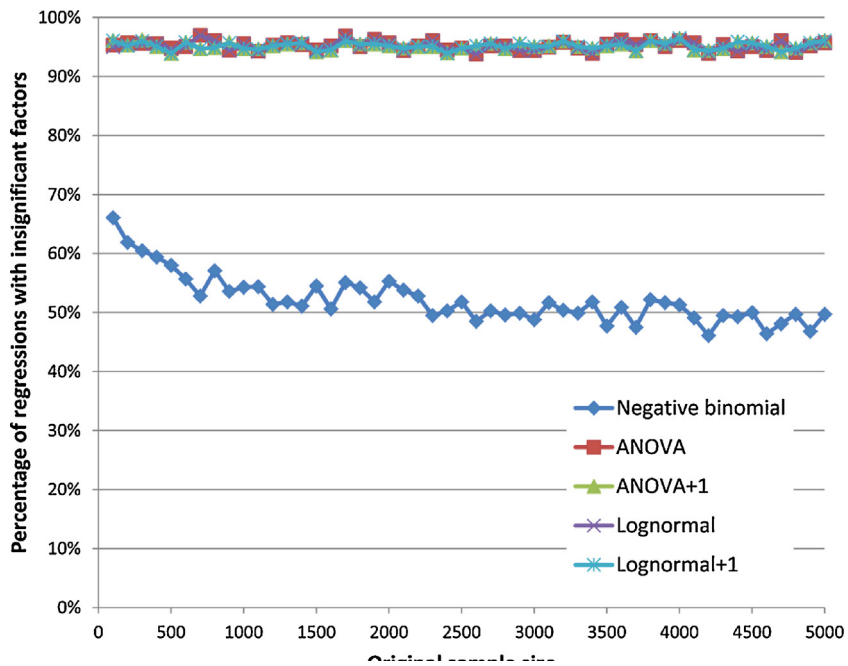


Fig. 2. Regressions against simulated discrete lognormal data (log mean: 0.5, log standard deviation: 2) at different sample sizes without factors ($n = 1000$ iterations). All lines overlap except for the negative binomial line.

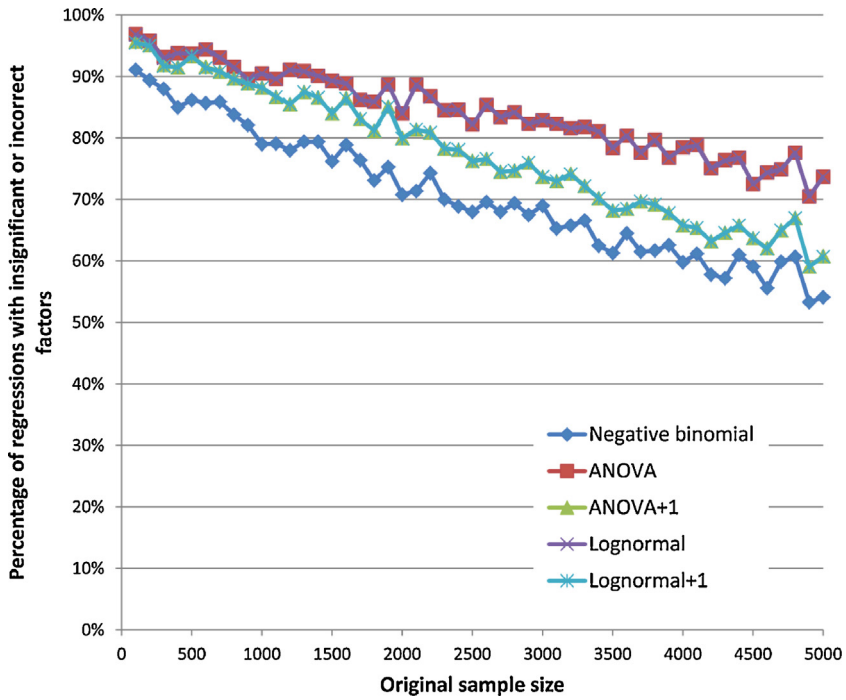


Fig. 3. Regressions against simulated discrete lognormal data (log standard deviation: 1) at different sample sizes with a binary factor with log means 0.5 and 0.55 ($n = 1000$ iterations). The ANOVA and lognormal lines and the lognormal + 1 and ANOVA + 1 lines overlap.

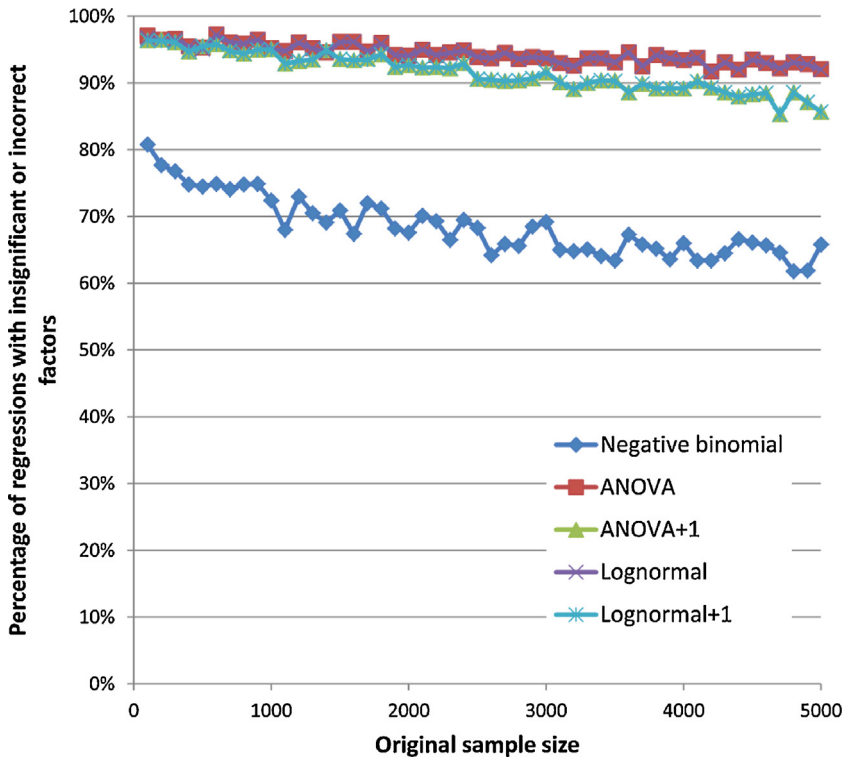


Fig. 4. Regressions against simulated discrete lognormal data (log standard deviation: 2) at different sample sizes with a binary factor with log means 0.5 and 0.55 ($n = 1000$ iterations). The ANOVA and lognormal lines and the lognormal + 1 and ANOVA + 1 lines overlap.

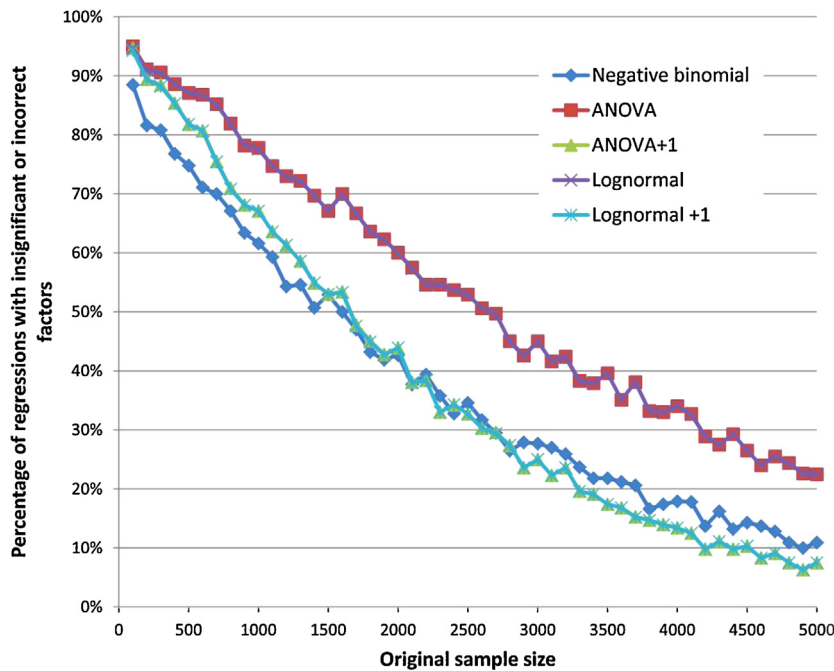


Fig. 5. Regressions against simulated discrete lognormal data (log standard deviation: 1) at different sample sizes with a binary factor with log means 0.5 and 0.6 ($n = 1000$ iterations). The ANOVA and lognormal lines and the lognormal+1 and ANOVA+1 lines overlap.

The results show that negative binomial regression applied to discrete lognormal data will identify non-existent factors in at a higher rate than expected by the significance level and so conclusions based upon negative binomial regression in such circumstances cannot be trusted. In contrast, the strategy of taking the logarithm of the lognormal data after discarding zeroes, and then applying the general linear model (ANOVA here, but other versions should also work) does not cause problems even though the resulting distribution is not normal. The situation for the generalised linear model with continuous lognormal error distribution is identical and the two approaches almost never give different results. The strategy of adding 1 to the data before taking its logarithm and then using the general linear model (as used by: Rigby, 2013; Tang, 2013) is similarly trustworthy and is equivalent to adding 1 to the data and then using the generalised linear model. The latter two are preferable, however, because they are more powerful in their ability to detect factors. This power to identify factors presumably derives from the larger sample size and increased difference between factors when the zeroes are not discarded.

In some cases, citation data has more uncited articles than would be expected by a standard statistical distribution (i.e., zero inflation). If this is suspected (e.g., by visual inspection of the distribution of citations for a particular dataset) then the zeroes can be removed and then either the log taken of the citations and the general linear model used (i.e., like the standard ANOVA model used above, without adding 1) or the generalised linear model used with the continuous lognormal distribution. Whilst these will be less powerful than using the full data set, they are unlikely to increase the chance of falsely identifying relationships and so are reasonable strategies.

If a discrete version of the lognormal distribution is created in the future for the generalised linear model then citation-like data could be regressed more precisely. If this happens then there would still be practical advantages with the above process (taking the log of the data and using the general linear model) because it seems intuitively unlikely that it would be substantially less powerful than the generalised linear model with the discrete lognormal distribution and it is a standard statistical test and hence there is fast and robust software for it that is widely understood and has many options.

The above strategies should be used for regressions of the mean (i.e., rather than logistic regression) for citation data that follows a discrete lognormal distribution. This seems likely to include most sets of articles with a common citation window but further studies are needed to identify general rules for when the lognormal distribution occurs. Nevertheless, researchers in the future should at least test for the lognormal distribution before using other models. The above strategies should also be considered for altmetrics, since they can correlate with citations (e.g., Li, Thelwall, & Giustini, 2012; Thelwall, Haustein, Larivière, & Sugimoto, 2013; Zahedi et al., in press) and hence it may have discrete lognormal distributions (e.g., article downloads: Yan & Gerstein, 2011), although preliminary studies of a range of altmetrics (data from: Thelwall et al., 2013) suggest that some may follow a power law instead.² Hence, the exact distributions of each one should be checked in case they follow different patterns.

² The full code for the tests, together with all the data used and additional graphs (including any marked as “not shown” in the text) is available at http://figshare.com/articles/Data_from_the_paper_Regression_for_Citation_Data/1090729.

References

- Abramo, G., Cicero, T., & D'Angelo, C. A. (2013). National peer-review research assessment exercises for the hard sciences can be a complete waste of money: The Italian case. *Scientometrics*, 95(1), 311–324.
- Abramo, G., D'Angelo, C. A., & Di Costa, F. (2011). National research assessment exercises: A comparison of peer review and bibliometrics rankings. *Scientometrics*, 89(3), 929–941.
- ACUMEN. (2014). *Guidelines for good evaluation practice with the ACUMEN Portfolio*. <http://research-acumen.eu/wp-content/uploads/Guidelines-Portfolio-TOC-document.v13dx.pdf>
- Adie, E., & Roe, W. (2013). Altmetric: Enriching scholarly content with article-level discussion and metrics. *Learned Publishing*, 26(1), 11–17.
- Aitchison, J., & Ho, C. H. (1989). The multivariate Poisson-log normal distribution. *Biometrika*, 76(4), 643–653.
- Aksnes, D. W. (2003). Characteristics of highly cited papers. *Research Evaluation*, 12(3), 159–170.
- Aksnes, D. W., Rørstad, K., Piro, F. N., & Sivertsen, G. (2013). Are mobile researchers more productive and cited than non-mobile researchers? A large-scale study of Norwegian scientists. *Research Evaluation*, 22(4), 215–223.
- Almind, T. C., & Ingwersen, P. (1997). Informetric analyses on the world wide web: Methodological approaches to 'webometrics'. *Journal of Documentation*, 53(4), 404–426.
- Anderson, D. L., Smart, W., & Tressler, J. (2013). Evaluating research – Peer review team assessment and journal based bibliographic measures: New Zealand PBRF research output scores in 2006. *New Zealand Economic Papers*, 47(2), 140–157.
- ARC. (2014). *Excellence in Research for Australia (ERA)*. <http://www.arc.gov.au/era/>
- Baldi, S. (1998). Normative versus social constructivist processes in the allocation of citations: A network-analytic model. *American Sociological Review*, 63(6), 829–846.
- Baumgartner, S. E., & Leydesdorff, L. (2014). Group-based trajectory modeling (GBTM) of citations in scholarly literature: Dynamic qualities of transient and sticky knowledge claims. *Journal of the Association for Information Science and Technology*, 65(4), 797–811.
- Borgman, C., & Furner, J. (2002). *Scholarly communication and bibliometrics. Annual review of information science and technology* (Vol. 36) Medford, NJ: Information Today, Inc.
- Bornmann, L., & Williams, R. (2013). How to calculate the practical significance of citation impact differences? An empirical example from evaluative institutional bibliometrics using adjusted predictions and marginal effects. *Journal of Informetrics*, 7(2), 562–574.
- Butler, L. (2003). Explaining Australia's increased share of ISI publications – The effects of a funding formula based on publication counts. *Research Policy*, 32(1), 143–155.
- Butler, L., & Mcallister, I. (2011). Evaluating university research performance using metrics. *European Political Science*, 10(1), 44–58.
- Chen, C. (2012). Predictive effects of structural variation on citation counts. *Journal of the American Society for Information Science and Technology*, 63(3), 431–449.
- Cole, J. R. (2000). *A short history of the use of citations as a measure of the impact of scientific and scholarly work. The Web of Knowledge: A Festschrift in Honor of Eugene Garfield*. Medford, NJ: Information Today.
- de Solla Price, D. (1965). Networks of scientific papers. *Science*, 169, 510–515.
- Didegah, F., & Thelwall, M. (2013a). Determinants of research citation impact in nanoscience and nanotechnology. *Journal of the American Society for Information Science and Technology*, 64(5), 1055–1064.
- Didegah, F., & Thelwall, M. (2013b). Which factors help authors produce the highest impact research? Collaboration, journal and document properties. *Journal of Informetrics*, 7(4), 861–873.
- Dobson, A. J., & Barnett, A. (2008). *An introduction to generalized linear models* (3rd ed.). Boca Raton, FL: Chapman & Hall/CRC.
- Dragos, C. M., & Dragos, S. L. (2014). Scientific productivity versus efficiency of R&D financing: Bibliometric analysis of African countries. *Current Science*, 106(7), 942–945.
- Evans, T. S., Hopkins, N., & Kaube, B. S. (2012). Universality of performance indicators based on citation and reference counts. *Scientometrics*, 93(2), 473–495.
- Foo, J. Y. A., & Tan, X. J. A. (2014). Analysis and implications of retraction period and coauthorship of fraudulent publications. *Accountability in Research*, 21(3), 198–210.
- Franceschet, M., & Costantini, A. (2011). The first Italian research assessment exercise: A bibliometric perspective. *Journal of Informetrics*, 5(2), 275–291.
- Garson, D. G. (2012). *General linear models: Univariate GLM, ANOVA/ANCOVA, repeated measures*. Asheboro, NC: Statistical Associates Publishers.
- Golosovsky, M., & Solomon, S. (2012). Runaway events dominate the heavy tail of citation distributions. *The European Physical Journal – Special Topics*, 205(1), 303–311.
- He, Z. L. (2009). International collaboration does not have greater epistemic authority. *Journal of the American Society for Information Science and Technology*, 60(10), 2151–2164.
- Hilbe, J. M. (2011). *Negative binomial regression* (2nd ed.). Cambridge: Cambridge University Press.
- Kousha, K., & Thelwall, M. (2014). Web impact metrics for research assessment. In B. Cronin, & C. Sugimoto (Eds.), *Beyond bibliometrics: Harnessing multidimensional indicators of scholarly impact* (pp. 289–306). Cambridge, MA: MIT Press.
- Kutlar, A., Kabasakal, A., & Ekici, M. S. (2013). Contributions of Turkish academicians supervising PhD dissertations and their universities to economics: An evaluation of the 1990–2011 period. *Scientometrics*, 97(3), 639–658.
- Leydesdorff, L., & Benschman, S. (2006). Classification and powerlaws: The logarithmic transformation. *Journal of the American Society for Information Science and Technology*, 57(11), 1470–1486.
- Li, X., Thelwall, M., & Giustini, D. (2012). Validating online reference managers for scholarly impact measurement. *Scientometrics*, 91(2), 461–471.
- MacRoberts, M. H., & MacRoberts, B. R. (1996). Problems of citation analysis. *Scientometrics*, 36(3), 435–444.
- Mavros, M. N., Bardakas, V., Rafailidis, P. I., Sardi, T. A., Demetriou, E., & Falagas, M. E. (2013). Comparison of number of citations to full original articles versus brief reports. *Scientometrics*, 94(1), 203–206.
- McCain, K. W. (2012). Assessing obliteration by incorporation: Issues and caveats. *Journal of the American Society for Information Science and Technology*, 63(11), 2129–2139.
- McDonald, J. D. (2007). Understanding journal usage: A statistical analysis of citation and use. *Journal of the American Society for Information Science and Technology*, 58(1), 39–50.
- Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63.
- Merton, R. K. (1973). *The sociology of science: Theoretical and empirical investigations*. Chicago: University of Chicago Press.
- Millar, R. B. (2011). *Maximum likelihood estimation and inference: With examples in R, SAS and ADMB*. New York: John Wiley & Sons.
- Mohammadi, E., & Thelwall, M. (2013). Assessing non-standard article impact using F1000 labels. *Scientometrics*, 97(2), 383–395.
- Mohammadi, E., & Thelwall, M. (2014). Mendeley readership altmetrics for the social sciences and humanities: Research evaluation and knowledge flows. *Journal of the Association for Information Science and Technology*, 65(8), 1627–1638.
- Mryglod, O., Kenna, R., Holovatch, Y., & Berche, B. (2013). Comparison of a citation-based indicator and peer review for absolute and specific measures of research-group excellence. *Scientometrics*, 97(3), 767–777.
- Peters, H. P. F., & van Raan, A. F. (1994). On determinants of citation scores: A case study in chemical engineering. *Journal of the American Society for Information Science*, 45(1), 39–49.
- Piwowar, H., & Priem, J. (2013). The power of altmetrics on a CV. *Bulletin of the American Society for Information Science and Technology*, 39(4), 10–13.
- Priem, J. (2014). Altmetrics. In B. Cronin, & C. Sugimoto (Eds.), *Beyond bibliometrics: Harnessing multidimensional indicators of scholarly impact* (pp. 263–287). Cambridge, MA: MIT Press.
- Priem, J., Taraborelli, D., Groth, P., & Neylon, C. (2010). *Altmetrics: A manifesto*. <http://altmetrics.org>

- Radicchi, F., Fortunato, S., & Castellano, C. (2008). Universality of citation distributions: Toward an objective measure of scientific impact. *Proceedings of the National Academy of Sciences*, 105(45), 17268–17272.
- REF. (2013). *Citation data*. <http://www.ref.ac.uk/subguide/citationdata/>
- Rigby, J. (2013). Looking for the impact of peer review: Does count of funding acknowledgements really predict research impact? *Scientometrics*, 94(1), 57–73.
- Seglen, P. O. (1992). The skewness of science. *Journal of the American Society for Information Science*, 43(9), 628–638.
- Shema, H., Bar-Ilan, J., & Thelwall, M. (2014). Do blog citations correlate with a higher number of future citations? Research blogs as a potential source for alternative metrics. *Journal of the Association for Information Science and Technology*, 65(5), 1018–1027.
- Sin, S. C. J. (2011). International coauthorship and citation impact: A bibliometric study of six LIS journals, 1980–2008. *Journal of the American Society for Information Science and Technology*, 62(9), 1770–1783.
- Sooryamoorthy, R. (2009). Do types of collaboration change citation? Collaboration and citation patterns of South African science publications. *Scientometrics*, 81(1), 177–193.
- Stasinopoulos, D. M., & Rigby, R. A. (2007). Generalized additive models for location scale and shape (GAMLSS) in R. *Journal of Statistical Software*, 23(7), 1–46.
- Stewart, J. A. (1983). Achievement and ascriptive processes in the recognition of scientific articles. *Social Forces*, 62(1), 166–189.
- Stringer, M. J., Sales-Pardo, M., & Amaral, L. A. N. (2010). Statistical validation of a global model for the distribution of the ultimate number of citations accrued by papers published in a scientific journal. *Journal of the American Society for Information Science and Technology*, 61(7), 1377–1385.
- Tabachnick, B. G., & Fidell, L. S. (2001). *Using multivariate statistics* (4th ed.). Needham Heights, MA: Allyn & Bacon.
- Tang, L. (2013). Does birds of a feather flock together matter—Evidence from a longitudinal study on US–China scientific collaboration. *Journal of Informetrics*, 7(2), 330–344.
- Thelwall, M., & Mafrahi, N. (2014). Are scholarly articles disproportionately read in their own country? An analysis of Mendeley readers. *Journal of the Association for Information Science and Technology* (in press).
- Thelwall, M., & Wilson, P. (2014). Distributions for cited articles from individual subjects and years. *Journal of Informetrics*, 8(4), 824–839. <http://www.scit.wlv.ac.uk/~cm1993/papers/DistributionsForCitedArticlesPreprintx.pdf>
- Thelwall, M., Haustein, S., Larivière, V., & Sugimoto, C. R. (2013). Do altmetrics work? Twitter and ten other social web services. *PLOS ONE*, 8(5), e64841.
- van Raan, A. F. J. (1998). The influence of international collaboration on the impact of research results. *Scientometrics*, 42(3), 423–428.
- Walters, G. D. (2006). Predicting subsequent citations to articles published in twelve crime–psychology journals: Author impact versus journal impact. *Scientometrics*, 69(3), 499–510.
- Waltman, L., van Eck, N. J., van Leeuwen, T. N., Visser, M. S., & van Raan, A. F. (2011). Towards a new crown indicator: Some theoretical considerations. *Journal of Informetrics*, 5(1), 37–47.
- Willis, D. L., Bahler, C. D., Neuberger, M. M., & Dahm, P. (2011). Predictors of citations in the urological literature. *BJU International*, 107(12), 1876–1880.
- Wu, J. (2013). Investigating the universal distributions of normalized indicators and developing field-independent index. *Journal of Informetrics*, 7(1), 63–71.
- Xia, J., & Nakanishi, K. (2012). Self-selection and the citation advantage of open access articles. *Online Information Review*, 36(1), 40–51.
- Yan, K. K., & Gerstein, M. (2011). The spread of scientific information: Insights from the web usage statistics in PLoS article-level metrics. *PLoS ONE*, 6(5), e19917.
- Yao, Z., Peng, X. L., Zhang, L. J., & Xu, X. J. (2014). Modeling nonuniversal citation distributions: The role of scientific journals. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(4), P04029.
- Yitzhaki, M. (1998). The 'language preference' in sociology: Measures of 'language self-citation', 'relative own-language preference indicator', and 'mutual use of languages'. *Scientometrics*, 41(1–2), 243–254.
- Yoshikane, F. (2013). Multiple regression analysis of a patent's citation frequency and quantitative characteristics: The case of Japanese patents. *Scientometrics*, 96(1), 365–379.
- Yu, T., Yu, G., & Wang, M. Y. (2014). Classification method for detecting coercive self-citation in journals. *Journal of Informetrics*, 8(1), 123–135.
- Zahedi, Z., Costas, R., & Wouters, P. (2014). How well developed are altmetrics? A cross-disciplinary analysis of the presence of 'alternative metrics' in scientific publications. *Scientometrics*, 101(2), 1491–1513.