

The role of statistics in establishing the similarity of citation distributions in a static and a dynamic context

Javier Ruiz-Castillo

Received: 17 August 2012 / Published online: 29 January 2013
© Akadémiai Kiadó, Budapest, Hungary 2013

Abstract Certain key questions in *Scientometrics* can only be answered by following a statistical approach. This paper illustrates this point for the following question: how similar are citation distributions with a fixed, common citation window for every science in a static context, and how similar are they when the citation process of a given cohort of papers is modeled in a dynamic context?

Keywords Citation distributions · Characteristic scores and scales · Power law distribution · Lognormal distribution · Log-location-scale family of distributions · Panel models · Multivariate models

Introduction

Consider the following basic question in *Scientometrics*: are citation distributions for different sciences very similar or rather different? This question is important for at least two reasons. Firstly, a deep understanding of the citation impact phenomenon requires explanatory models in which thousands of researchers decide in a decentralized manner which of their colleagues' previous works should be cited. Consider the following three possibilities: (i) all distributions are essentially different; (ii) the distributions for the natural sciences are different from the distributions for the social sciences, and (iii) all distributions are broadly similar. Whichever of these turns out to be the case may drastically affect the type—or even the number—of explanatory models we need. Secondly, we know that publication and citation practices are very different across scientific fields. This makes the direct comparison of citation counts of articles belonging to different fields extraordinarily difficult. However, the closer citation distributions are to each other, the

J. Ruiz-Castillo (✉)
Department of Economics, Universidad Carlos III, Getafe, Madrid, Spain
e-mail: jrc@eco.uc3m.es

easier it might be to find sensible normalization procedures to compare citation counts across fields, and to aggregate a set of heterogeneous sub-fields into larger categories.¹

We should distinguish between two scenarios. In the first one, citation windows are fixed and equal for all sciences. In this static case, there is systematic evidence showing that, to a first approximation, citation distributions are strikingly similar at different aggregation levels. Moreover, confirming all previous, casual but valuable evidence, we do know that citation distributions are highly skewed (Price 1965; Seglen 1992; Schubert and Braun 1986; Glänzel 2007; Albarrán and Ruiz-Castillo 2011; Albarrán et al. 2011). Such similarity has recently lead some authors towards a universality claim according to which, except for a scale factor, citation distributions are equally distributed (Radicchi et al. 2008; Radicchi and Castellano 2012a, b). There is evidence that the claim does not hold at the lower or the upper tail of citation distributions. Nevertheless, it appears that there is “enough” similarity among citation distributions to make meaningful comparisons of citation counts across different fields, and to justify simple but practical normalization procedures for aggregating heterogeneous sub-fields or disciplines into larger scientific categories (Glänzel 2011; Radicchi and Castellano 2012a; Crespo et al. 2012).

The second scenario focus on the dynamic citation process of a given set of papers published in the same period. There are two types of interesting questions to study. Firstly, one may simply ask: how long does it take for citation distributions in different sciences to adopt their characteristic skewed shape? (Glänzel 2007). Secondly, one may build an explanatory model for the probability that each article will receive one more citation as a function of a number of factors including past citations, the characteristics of the article, its author(s), and the field to which the article belongs. In a multivariate context, this question remains essentially unexplored.

These are old issues, pregnant with methodological, substantive, and policy consequences that have been fruitfully discussed in *Scientometrics* for several decades using a variety of approaches. The point of this paper is that both specialists in *Scientometrics* and scientists of all persuasions will not accept a definite answer that is not firmly rooted in rigorous statistical methods. This is, of course, why I suggest reviewing these questions for this monographic issue of *Scientometrics*. In particular, given the wide differences in mean citation and field size, to study the skewness of science it is convenient to use scale- and size-independent techniques; ideally, to establish the universality claim we need a methodology that consists of several well defined stages, and to study the probability that each article will receive one more citation dynamic, multivariate models are strongly recommended.

The rest of the paper is organized into four sections. “[The skewness of science](#)” and “[The universality claim](#)” sections briefly reviews the skewness of science and the universality claim in the static case. “[Dynamic considerations](#)” section discusses the characteristics of a dynamic model for the probability that an article gets one more citation, while “[Conclusions](#)” section offers some brief concluding comments.

The skewness of science

To begin with, we would like to know whether—to a first approximation—citation distributions share or not a set of basic features. In particular, we would like to know whether common characteristics are only present at a high aggregation level and disappear at the

¹ For reasons of space, only methods that belong to the class of target or “cited side” normalization procedures will be discussed in this paper.

sub-field level or, on the contrary, whether they are present at the lowest aggregation level and, in this case, whether they are preserved or not at higher aggregate levels.

To answer this question, sub-fields are identified with the 219 Web of Science subject-categories distinguished by Thomson Scientific. Articles are assigned to sub-fields through the journals in which they have been published. The problem is that many journals are assigned to a single sub-field, but many others are assigned to multiple sub-fields. To handle this difficulty, there are two alternatives: a *fractional* strategy, according to which each publication is fractioned into as many equal pieces as necessary, with each piece assigned to a corresponding sub-field, and a *multiplicative* strategy according to which each paper is wholly counted as many times as necessary in the several sub-fields to which it is assigned.

As indicated in the “**Introduction**” section, size- and scale-independent descriptive tools permit us to focus on the *shape* of these 219 distributions. In particular, the characteristic scores and scales approach, pioneered by Schubert and Braun (1986) in *Scientometrics*, permits the partition of any distribution of articles into classes according to the citations they receive. Let s_1 denote the mean citation, and s_2 the mean citation of articles above s_1 . Consider the partition of citation distributions into three broad classes of articles that (I) receive no or few citations below s_1 , (II) are fairly well cited, namely, with citations between s_1 and s_2 , and (III) are remarkably or outstandingly cited with citations above s_2 . Using a large dataset consisting of about 3.7 million articles published in 1998–2002 with a five-year citation window, Albarrán et al. (2011) partition every sub-field citation distribution into these three classes, and then compute the percentage of total citations corresponding to each class. For both the multiplicative and the fractional strategies, Table 1 provides information about the average percentage of articles (and standard deviation) in classes I and III over the 219 distributions, as well as the percentage of total citations in each class.

As expected, citation distributions are highly skewed in the sense that, on average, 69 % of all articles receive citations below the mean and account for 21 % of all citations, while articles with a remarkable or outstanding number of citations represent about 9 or 10 % of the total, and account for approximately 45 % of all citations. We should emphasize that the small standard deviations in Table 1 indicate that this description provides a good characterization of citation distributions at the sub-field level. Since sub-field shapes are so similar, any reasonable aggregation scheme should preserve its main features. This is exactly what is found when sub-fields are aggregated into larger categories (For the multiplicative and the fractional cases, see Albarrán et al. (2011) and Herranz and Ruiz-Castillo (2012), respectively).²

It can be concluded that the celebrated title of Seglen’s (1992) seminal contribution appropriately summarizes the massive available evidence about citation distributions at several aggregation levels. This is a very important result, pointing towards a single theoretical explanation of this aspect of the scientific communication process. Who could have thought so, given the wide differences in professional culture, publication and citation practices across scientific fields, let alone the myriad of motives researchers may have to cite their colleagues’ previous work?

² Of course, from the beginning of *Scientometrics* it is generally believed that citation distributions are highly skewed (Price 1965; Seglen 1992). However, until the results previously summarized, the empirical evidence, although valuable, was very scant and mostly unsystematic. See Albarrán et al. (2011) for a discussion of the previous literature, including the limitations of the few systematic studies by Schubert and Braun (1986), Glänzel (2007) and Albarrán and Ruiz-Castillo (2011).

Table 1 Characteristic scores and scales

	Percentage of articles in classes			
	I	III	I	III
Multiplicative strategy	68.6 (3.7)	10.0 (1.7)	21.1 (5.0)	44.9 (4.6)
Fractional strategy	68.3 (3.4)	10.2 (1.6)	21.5 (4.2)	44.7 (3.9)

Average values (and standard deviations) over 219 sub-fields in the partition of citation distributions into three broad categories

The universality claim

The similarity just discussed opens the way for the next question: independently of their means and sizes, are citation distributions equally distributed? In practice, we rarely, if ever, know for certain that an observed quantity is drawn from a given distribution. The best we can typically do is to say that our observations are consistent with a model of the world in which the quantity of interest—the number of citations received by an article, in our case—follows a certain distribution. Unfortunately, there is no easy short cut to reach conclusions of this sort. The standard process for estimating and testing which is the best distribution that fits some data can be summarized in the following three steps. Firstly, under the assumption that the data are drawn from a distribution of known functional form, one must first find the best available estimates and standard errors of the parameters of the distribution in question. Secondly, using an appropriate measure of the goodness of fit we must decide whether the previous distributional assumption is even a reasonable hypothesis to begin with. Thirdly, the fact that a distribution cannot be rejected does not guarantee that it is the best of all possibilities. New tests have to be applied to know which of, say, two alternative models is the better fit to the data (see Clauset et al. 2009, for a description of the process in the power law case).

There are three major, systematic studies that have followed the first two steps in this procedure (except for the estimation of the parameter standard errors). Firstly, Albarrán et al. (2011) find that the existence of a power law cannot be rejected at the top of the upper tail in 140, or 64 % of 219 sub-fields. On the one hand, it is really impressive that we cannot reject the hypothesis of a power law in so many cases. However, this result shows that the existence of a power law is a prevalent but not a universal phenomenon. It should be mentioned that, when they cannot be rejected, power laws only represent a small proportion of the upper tail of citation distributions but account for a considerable percentage of all citations: on average, power laws represent 2 % of all articles in a sub-field, and account for about 13.5 % of all citations. Finally, even in the absence of confidence intervals, the large dispersion of the power law parameters (see Panel B in Table 2 in Albarrán et al. (2011)) is a clear indication that excellence is not equally structured in all citation distributions for which a power law cannot be rejected.³ Secondly, using other methods, Waltman and van Eck (2012) reach the same conclusion.

Thirdly, using a dataset of about three million papers, covering 172 subject-categories, Radicchi and Castellano (2012a)—RC hereafter—conclude that the log normality of

³ In addition, consider the possibility of defining a high-impact indicator over the sub-set of articles with citations above a high percentile of citation distributions, say the 80th percentile. The distribution of high-impact values for the 219 sub-fields according to an indicator of this type is highly skewed to the right, and it presents some important extreme observations (see Herranz and Ruiz-Castillo 2012).

citation distributions cannot be rejected for a very limited number of subject-categories: 37 % of the sub-fields for papers published in 1980, 28 % in 1985, 20 % in 1990, 10 % in 1995, 5 % in 1999, and 4 % in 2004. From this and other evidence, RC conclude that differences between citation distributions cannot only be due to a single scaling factor. Therefore, using a linear transformation to normalize citation distribution does not seem justified. However, the usual practice in *Scientometrics* (see inter alia Moed et al. 1985, 1995; Moed and van Raan 1988; Braun et al. 1985; Schubert et al. 1983, 1987a, b, 1988; Schubert and Braun 1986, 1996; Vinkler 1986, 2003), advocated by these authors in other contributions (Radicchi et al. 2008; Radicchi and Castellano 2012b), consists in dividing the raw number of citations by a single discipline dependent scaling factor: the sub-field mean citation.

In order to make further progress in the normalization issue, RC use a sort of “reverse engineering” approach that starts by choosing a reference distribution: the cumulative distribution of raw citation counts obtained by aggregating together their 172 Web of Science categories in the multiplicative case. Then one searches for a transformation of the raw data in every sub-field, so that each distribution of transformed citation counts coincides with the reference distribution. Very interestingly, with few exceptions⁴ the transformation for each sub-field s can be described by a non-linear relationship depending on two parameters, the pre-factor, a_s , and the exponent of the mapping, α_s :

$$\text{Transformed citation count of an article in sub-field } s = (\text{Raw citation count}/a_s)^{1/\alpha_s}. \quad (1)$$

Three aspects of this research should be treated separately. Firstly, in general the transformation effectively removes systematic differences in citation practices among sub-fields in the top 10 % of highly cited papers (see Fig. 4 in Radicchi and Castellano (2012a), for papers published in 1999). This is an extremely useful result that constitutes the best available justification for a specific normalization procedure depending on only two parameters. Secondly, RC claim that, in spite of the exceptions already mentioned, their results support the claim that citation distributions are universal, in the sense that they are all part of the same family of univariate distributions, namely, a *log-location-scale family* of distributions that can be rewritten in terms of the same reference distribution for any choice of a location and a scale parameter (Weinbull and log-normal distributions are well known log-location-scale families). In our opinion, this claim should be qualified. RC’s results provide estimates of the parameters of a log-location-scale distribution for 172 sub-fields. We would like to see the standard errors of the estimated parameters, results on the goodness of fit, and the testing against alternative models. Thirdly, in the absence of conclusive statistical evidence about confidence intervals, RC claim not only that the parameters are very stable across sub-fields and publication years, but that the strong regularities in the size of the exponent of the mapping make the transformation almost linear for the majority of sub-fields. Hence, the usual practice of rescaling raw citation counts by the sub-field mean, despite not being strictly correct, can be considered a very good approximation of the transformation making citation counts independent on the scientific domain. In our opinion, more statistical work should be done before definitively accepting these exciting conclusions.

⁴ In general, these exceptions are hybrid subject categories, like the “Multiplidisciplinary sciences”, or not so well defined sub-fields, like “Engineering, petroleum” or “Biodiversity conservation”, whose papers are also assigned, through the journals where they have been published, to other broader subject-categories.

Dynamic considerations

So far, we have restricted the discussion to the static case in which citation windows are fixed and equal for all sciences. However, the question of the similarity of citation distributions cannot be finished without reference to a situation in which citations and other characteristics of a given set of articles published in a certain period are observed during a number of time periods. Two comments are in order.

Firstly, any sensible model for the probability that an article will receive one more citation must include a variety of factors of two types. Some factors might be time invariant, such as (i) what we may call the article's a priori scientific influence; (ii) other characteristics of the article, such as the number of authors, the degree of international co-authorship, or the journal in which it has been published; (iii) characteristics of the author(s), including their past citation impact, and other exogenous variables like their age or gender, or (iv) the characteristics of the reward system in science where material resources and recognition are allocated on the basis of past performance. However, other influential factors may vary with time, such as (v) the sequence of citations received in previous periods; (vi) characteristics of the field, such as its size or the percentage of references that become citations in the periodical literature, or simply (vii) luck. Secondly, the mere description of the question indicates that our multivariate data has a panel or longitudinal structure. The inclusion of time invariant variables permit to improve our understanding of the dynamic aspects of the problem controlling for fixed effects. On the other hand, because panel data have both cross-sectional and time series dimensions, the application of regression models to fit econometric models are more complex than those for simple cross-sectional data sets. However, there are suitable methods for longitudinal data developed in Statistics and Econometrics, such as panel data regression methods (Arellano and Honoré 2001; Arellano and Bonhomme 2011), or mixed models that use parametric fixed effects to represent the covariate effects, and an arbitrary smooth function to model the time effect (Diggle et al. 2002; Molenberghs and Verbeke 2005).

Nothing is new in the above partial listing of potential explanatory variables. What is new is the framework itself, which allows certain classical issues to be posed in a new light. We would mention four examples. Firstly, consider the well-known *Matthew effect* due to Merton (1968, 1988), according to which small differences in initial status amplify over time to generate cumulative advantages—a process also known as “preferential attachment” in network analysis. Many papers have discussed and have offered some empirical evidence about the presence of a Matthew effect explaining some of the success of highly cited papers. But very few have seriously controlled for confounding factors in a multivariate framework, beginning with the scientific influence attributable to different articles and/or the talent for citation impact attributable to different authors. For reasons of space, I can only cite the remarkable contribution of Azoulay et al. (2012), which cleverly designs a dataset for the estimation of the Matthew effect controlling for the characteristics of both the authors and the articles involved. Specifically, they estimate the effect of a scientist becoming a Howard Hughes medical investigator on citations to articles published *before* the prize was awarded. They find evidence of a small post-appointment citation boost confined to a short time window. However, the effect of the prize is significantly larger when there is uncertainty about producer and product quality; for instance, the consequence of a shift in status is much larger for younger producers, for those whose positions in the status ordering are less well established at the time they receive the prize, for papers published in lower quality journals, and for articles written in novel areas of science.

Secondly, using a panel data set on individual researchers from the University of Leuven, comprising 10 years of publication data 1992–2001, Kelchtermans and Veugelers (2012) study top research performance and its persistence over time. They study the factors that influence the hazard for a researcher to achieve a first and subsequent top performance level, taking into account time-varying and invariant covariates, and checking for the influence of past (top) performance. The results confirm the skewness of research productivity, as well as the persistency of top performance. On an important topic that has been worrying researchers since Cole and Zuckerman (1984), the results consistently show the importance of gender as a determining factor, with females being significantly less likely to reach first top performance. However, once women break through to their first top performance, no gender bias hinders them in future top performance. On the contrary, women seem to persist in top performance more easily than men. Interestingly, all the results are robust to corrections for unobservable individual heterogeneity, suggesting that factors like “talent” remain an integral part of the story of top performance and its persistence. Finally, the substantial part of the remaining residual left unexplained indicates that luck cannot be excluded from the story either. The richness of the results shows the dramatic advantages of working with multivariate statistical models.

Thirdly, Katz (2000) presents some interesting evidence on the (unconditional) relationship between the citations received by articles in a given field and the size of the field in question measured by the total number of articles. Clearly, this important question should be studied in a multivariate context where the effect of many relevant variables is controlled for. Fourthly, last but not least, a model of this type will allow us to obtain conclusions about the main topic discussed in this paper—how similar citation distributions are—taking into account the effect of many variables.

Conclusions

As Clauset et al. (2009) indicate, “*Scientists have learned many things from observation of the statistical distributions of experimental quantities such as the lifetimes of excited atomic or particle states, populations of animals, plants, or bacteria, prices in the stock market, or the arrival times of messages sent across the Internet.*” Similarly, whether citation distributions are very similar or rather different is a basic question in *Scientometrics* that can only be answered in statistical terms. For a common, fixed citation window in the static case the empirical conclusions that can provisionally drawn are the following two.

1. Citation distributions at the lowest level of aggregation are remarkably similar. This has already opened the way towards simple but effective normalization procedures to aggregate heterogeneous sub-fields into larger categories.
2. Whether it can be claimed that all (or most) citation distributions (i) belong to the same family, and (ii) except for a scale parameter, all other parameters are statistically equal, is still an open question. Four comments should be emphasized. Firstly, until now systematic studies at the lowest aggregation level have taken as given the list of Web of Science categories distinguished by Thomson Scientific—a classification scheme in need of serious revision. Thus, part of the difficulties encountered by formal procedures may disappear whenever a better classification scheme becomes available. Secondly, even if the zero citation phenomenon at a comparable stage of the citation process were to be the same in all sciences, when we take a fixed, common citation

window for all of them we are bound to have very different percentages of articles without citations: a small (large) percentage in quick (slow) sciences.⁵ In so far as the citation process evolves at a different velocity in different fields, using variable citation windows to ensure that the process has reached a similar stage in all fields should improve field comparability at the lower tail of citation distributions. Thirdly, as RC point out, the discreteness of citation numbers becomes more relevant at the lower tail of citation distributions, making it harder for any normalization procedure to work well in this segment. Fourthly, Crespo et al. (2012) study the evolution of citation inequality attributable to differences in citation practices at different quantiles of citation distributions. In the case of the 22 broad fields distinguished by Thomson Scientific, they find that this type of citation inequality is particularly high until, approximately, the 70th percentile, as well as for a few quantiles at the very upper tail of citation distributions; otherwise, this citation inequality appears to be approximately constant for a wide range of intermediate values. The conjecture is that citation distributions are truly close to each other except for a scale factor only in such an intermediate range of values—a fact that might be taken into account in further research for a universal characterization of citation distributions.

The similarity between citation distributions already discussed seems to imply that all we need is a single behavioral explanation of this aspect of the scientific communication process. However, before this task is finally attacked, it would be convenient to know more about the similarities—or differences—of the evolution of citation distributions for a given set of articles as the citation window varies over time. Dynamic, multivariate models constitute a promising statistical tool for approaching this topic.

Acknowledgments Financial support from the Spanish MEC, through Grant No. SEJ2007-67436, as well as conversations with Pedro Albarrán are gratefully acknowledged. A referee report helped to improve the original version of the paper. All shortcomings are the author's sole responsibility.

References

- Albarrán, P., Crespo, J., Ortuño, I., & Ruiz-Castillo, J. (2011). The skewness of science in 219 sub-fields and a number of aggregates. *Scientometrics*, 88, 385–397.
- Albarrán, P., & Ruiz-Castillo, J. (2011). References made and citations received by scientific articles. *Journal of the American Society for Information Science and Technology*, 62, 40–49.
- Arellano, M., & Bonhomme, S. (2011). Nonlinear panel data analysis. *Annual Review of Economics*, 3, 395–424.
- Arellano, M., & Honoré, B. (2001). Panel data models: Some recent developments. In J. Heckman & E. Leamer (Eds.), *Handbook of econometrics* (Vol. 5). Amsterdam: North Holland.
- Azoulay, P., Stuart, T., & Wang, Y. (2012). Mathew: Effect or Fable? NBER Working Paper No. 18625. (<http://www.nber.org/papers/w18625>).
- Braun, T., Glänzel, W., & Schubert, A. (1985). *Scientometrics indicators. A 32 country comparison of publication productivity and citation impact*. Singapore: World Scientific Publishing Co. Pty. Ltd.
- Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review*, 51, 661–703.

⁵ This is exactly what Albarrán et al. (2011) find with the dataset mentioned in “The skewness of science” section: over the 219 sub-fields in the multiplicative case, for example, the average percentage of un-cited articles is 24.7 % and the standard deviation 13.9. Therefore, the coefficient of variation is 0.56, while the one for the percentage of articles in class I in Table 1 is 0.05. On the other hand, RC also find that the transformation of Eq. 1 becomes less descriptive beyond the top 10 % of highly cited articles, and the removal of the bias in the raw data worsens.

- Cole, J. R., & Zuckerman, H. (1984). The productivity puzzle: Persistence and change in patterns of publication of men and women scientists. *Advances in Motivation and Achievement*, 2, 217–258.
- Crespo, J. A., Li, Y., & Ruiz-Castillo, J. (2012). The comparison of citation impact across scientific fields. Working paper 12-06. Madrid: Universidad Carlos III. (<http://hdl.handle.net/10016/14771>).
- Diggle, P. J., Heagerty, P., Liang, K.-Y., & Zeger, S. L. (2002). *Analysis of longitudinal data* (2nd ed.). Oxford: Clarendon Press.
- Glänzel, W. (2007). Characteristic scores and scales: A bibliometric analysis of subject characteristics based on longterm citation observation. *Journal of Informetrics*, 1(1), 92–102.
- Glänzel, W. (2011). The application of characteristic scores and scales to the evaluation and ranking of scientific journals. *Journal of Information Science*, 37, 40–48.
- Herranz, N., & Ruiz-Castillo, J. (2012). Multiplicative and fractional strategies when journals are assigned to several sub-fields. *Journal of the American Society for Information Science and Technology*. (doi: 10.1002/asi.22629); in press.
- Katz, J. S. (2000). Scale-independent indicators and research evaluation. *Science and Public Policy*, 27, 23–36.
- Kelchtermans, S., & Veugelers, R. (2012). Top research productivity and its persistence: Gender as a double-edged sword. *The Review of Economics & Statistics* (forthcoming).
- Merton, R. K. (1968). The Matthew effect in science. *Science*, 159, 56–63.
- Merton, R. K. (1988). The Matthew effect in science, II: Cumulative advantage and the symbolism of intellectual property. *Isis*, 79, 606–623.
- Moed, H. F., Burger, W. J., Frankfort, J. G., & van Raan, A. F. J. (1985). The use of bibliometric data for the measurement of university research performance. *Research Policy*, 14, 131–149.
- Moed, H. F., De Bruin, R. E., & van Leeuwen, T. N. (1995). New bibliometrics tools for the assessment of national research performance: Database description, overview of indicators, and first applications. *Scientometrics*, 33, 381–422.
- Moed, H. F., & van Raan, A. F. J. (1988). Indicators of research performance. In A. F. J. van Raan (Ed.), *Handbook of quantitative studies of science and technology* (pp. 177–192). Amsterdam: North Holland.
- Molenberghs, G., & Verbeke, G. (2005). *Models for discrete longitudinal data*. New York: Springer.
- Price, D. D. (1965). Networks of scientific papers. *Science*, 149, 510–515.
- Radicchi, F., & Castellano, C. (2012a). A reverse engineering approach to the suppression of citation biases reveals universal properties of citation distributions. *PLoS One*, 7(e33833), 1–7.
- Radicchi, F., & Castellano, C. (2012b). Testing the fairness of citation indicators for comparisons across scientific domains: The case of fractional citation counts. *Journal of Informetrics*, 6, 121–130.
- 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 of the United States of America*, 105, 17268–17272.
- Schubert, A., & Braun, T. (1986). Relative indicators and relational charts for comparative assessment of publication output and citation impact. *Scientometrics*, 9, 281–291.
- Schubert, A., & Braun, T. (1996). Cross-field normalization of scientometric indicators. *Scientometrics*, 36, 311–324.
- Schubert, A., Glänzel, W., & Braun, T. (1983). Relative citation rate: A new indicator for measuring the impact of publications. In D. Tomov & L. Dimitrova (Eds.), *Proceedings of the First National Conference with International Participation in Scientometrics and Linguistics of Scientific Text* (pp. 80–81). Varna, Bulgaria: ICS.
- Schubert, A., Glänzel, W., & Braun, T. (1987a). Subject field characteristic citation scores and scales for assessing research performance. *Scientometrics*, 12, 267–292.
- Schubert, A., Glänzel, W., & Braun, T. (1987b). Subject field characteristic citation scores and scales for assessing research performance. *Scientometrics*, 12, 267–292.
- Schubert, A., Glänzel, W., & Braun, T. (1988). Against absolute methods: Relative scientometric indicators and relational charts as evaluation tools. In A. F. J. van Raan (Ed.), *Handbook of quantitative studies of science and technology* (pp. 137–176). Amsterdam: Elsevier.
- Seglen, P. (1992). The skewness of science. *Journal of the American Society for Information Science*, 43, 628–638.
- Vinkler, P. (1986). Evaluation of some methods for the relative assessment of scientific publications. *Scientometrics*, 10, 157–177.
- Vinkler, P. (2003). Relations of relative scientometric indicators. *Scientometrics*, 58, 687–694.
- Waltman, L., & van Eck, N. J. (2012). Universality of citation distributions revisited. *Journal of the American Society for Information Science and Technology*, 63, 72–77.