



Predicting the long-term citation impact of recent publications[☆]



Clara Stegehuis^{a,*}, Nelly Litvak^b, Ludo Waltman^c

^a Eindhoven University of Technology, Department of Mathematics and Computer Science, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

^b University of Twente, Department of Applied Mathematics, P.O. Box 217, 7500 AE Enschede, The Netherlands

^c Centre for Science and Technology Studies, Leiden University, P.O. Box 905, 2300 AX Leiden, The Netherlands

ARTICLE INFO

Article history:

Received 30 March 2015

Received in revised form 14 June 2015

Accepted 15 June 2015

Keywords:

Citation analysis

Citation impact

Impact factor

Prediction

Quantile estimation

Quantile regression

ABSTRACT

A fundamental problem in citation analysis is the prediction of the long-term citation impact of recent publications. We propose a model to predict a probability distribution for the future number of citations of a publication. Two predictors are used: the impact factor of the journal in which a publication has appeared and the number of citations a publication has received one year after its appearance. The proposed model is based on quantile regression. We employ the model to predict the future number of citations of a large set of publications in the field of physics. Our analysis shows that both predictors (i.e., impact factor and early citations) contribute to the accurate prediction of long-term citation impact. We also analytically study the behavior of the quantile regression coefficients for high quantiles of the distribution of citations. This is done by linking the quantile regression approach to a quantile estimation technique from extreme value theory. Our work provides insight into the influence of the impact factor and early citations on the long-term citation impact of a publication, and it takes a step toward a methodology that can be used to assess research institutions based on their most recently published work.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Citation counts are a popular indicator of the impact of scientific publications. In the evaluation of research institutions, bibliometric indicators based on the citations received by the publications of an institution often play an important role. However, the use of citation-based indicators is problematic when the impact of recent publications needs to be determined. One or two years after their appearance, most publications have received only a few citations. After one year, there are many publications with just one or two citations or even with no citations at all. Some of these publications may receive a lot of citations in later years, while others may attract hardly any attention in the future. This makes it difficult to determine the impact of recent publications. Nevertheless, research institutions often want their performance to be assessed based on their most recent work (Bornmann, 2013). For example, a research institution may be interested in comparing the performance of different research groups based on publications from the four or five most recent years. Publications from earlier years may not be considered of interest anymore because many of them were produced by researchers who are no longer affiliated to

[☆] The peer review process of this paper was handled by Vincent Larivière, Associate Editor of Journal of Informetrics.

* Corresponding author.

E-mail address: c.stegehuis@tue.nl (C. Stegehuis).

the institution or because they belong to old research lines that in the meantime have been closed down. Given the practical importance of taking into account very recent publications in research evaluations, our aim in this paper is to introduce a model for making predictions of the impact that recent publications will have in the long term.

Our model predicts the long-term citation impact of a publication based on two variables, namely the impact factor of the journal in which the publication has appeared and the number of early citations the publication has received. Early citations are defined as citations received in the year in which a publication appeared or in the year thereafter. The two predictors that we use are easily available, and contrary to for example the prediction approach proposed by Wang, Song, and Barabási (2013), they allow predictions to be made fairly soon after the appearance of a publication. Also, compared with other predictors that could be considered, such as the length of the reference list of a publication or the number of authors of a publication, the predictors that we use are relatively hard to manipulate. Earlier studies have shown that both impact factors and early citations are important predictors of future citations. In the next section, we will provide an overview of these earlier studies and we will discuss their relationship with our present work.

Earlier studies on citation impact prediction have often focused on providing a point estimate of the future number of citations of a publication. Given the high degree of uncertainty in citation impact predictions, we believe that it is more relevant to know the probability that a publication will receive a certain number of citations in the future. We therefore do not predict the average number of citations that a publication is expected to attract in the future, but instead we predict a probability distribution for the future number of citations of a publication. To predict this probability distribution conditional on a publication's impact factor and its early citations, we employ the technique of quantile regression introduced by Koenker and Bassett (1978).

We also study the relationship between our prediction model based on quantile regression and results from extreme value theory. To do so, we first use so-called Zenga plots, introduced recently by Cirillo (2013), to establish that the citation distributions obtained in our analysis have a Pareto tail. This result then enables us to provide analytical insight into the behavior of the quantile regression coefficients for high quantiles. More specifically, we are able to link the regression coefficients to an estimator for the tail quantiles of a Pareto distribution developed in the framework of extreme value theory (Dekkers, Einmahl, & De Haan, 1989).

We use citation data for a large set of publications in the field of physics to test our prediction approach. The data is taken from the Web of Science database.

The paper is organized as follows. First, Section 2 discusses how our research relates to earlier work reported in the literature. Next, Section 3 describes the data that were used in our analysis. Section 4 then introduces our model for predicting the long-term citation impact of publications, conditional on impact factors and early citations. Section 5 presents our empirical results. Sections 5.1 and 5.2 focus on the values obtained for the parameters of our model. Sections 5.3–5.6 address the fit of this model to the data and the predictive power of the model. Section 6 studies the relationship between our model and results from extreme value theory. Section 7 addresses the sensitivity of the parameters of our model to differences between fields of science, focusing on the fields of biology, chemistry and physics. Finally, Section 8 concludes the paper.

2. Relation with earlier work

There is an extensive literature on modeling or predicting the number of citations of a publication based on all kinds of variables. An early study in this literature is the work by Peters and Van Raan (1994), who investigate the determinants of the citation impact of chemical engineering publications. More recent work in this literature is reported by, among others, Bornmann, Leydesdorff, and Wang (2013), Didegah and Thelwall (2013a, 2013b), Fu and Aliferis (2010), Haslam et al. (2008), Walters (2006), Wang et al. (2012), Wang, Yu, and Yu (2011), Yu, Yu, Li, and Wang (2014), and Onodera and Yoshikane (2015). Various studies have also appeared in non-bibliometric journals (e.g. Haslam & Koval, 2010; Lokker, McKibbin, McKinlay, Wilczynski, & Haynes, 2008; Mingers & Xu, 2010). Recent overviews of the literature on modeling or predicting citation impact are provided by Didegah and Thelwall (2013a, 2013b) and Onodera and Yoshikane (2015). Examples of variables that have been found to predict citation impact include the impact factor of the journal in which a publication has appeared, the type of study (e.g., original research vs. literature review), the number of pages of a publication, the number of references of a publication, the number of authors, institutions, and countries in a publication's address list, and the past performance of these authors, institutions, and countries.

It is important to emphasize that the objective of our work is different from the studies mentioned above. Like the above-mentioned studies, our interest is in predicting citation impact. However, our more specific interest is in using citation impact predictions in the evaluation of researchers, research groups, research institutions, and so on. In this specific context, many of the variables that have been found to correlate with citation impact should not be used for making citation impact predictions. Some variables have the problem that they can be easily manipulated. For example, suppose researchers know that they will be evaluated based on the predicted citation impact of their publications, and suppose researchers also know that the citation impact of a publication will be predicted based on, for example, the number of pages or the number of references of the publication. In that case, in order to be evaluated more favorably, it might be tempting for researchers to try to artificially increase the number of pages or the number of references of their publications. Therefore, we consider only variables that cannot be altered easily by the authors of a publication. Other variables have the problem that they may lead to an undesirable self-reinforcing effect. For example, suppose researchers are evaluated based on the predicted

citation impact of their publications, and suppose the citation impact of a publication is predicted based on the citation impact of the earlier work of the authors of the publication. In that case, researchers who were successful in their older work will automatically be predicted to be successful also in their more recent work. This creates a self-reinforcing effect. Future success is determined by past success.

In order to avoid problems related to manipulation and self-reinforcing effects, we aim to predict the citation impact of a publication based on indicators that are available shortly after the publication's appearance and that can be considered to provide an impression of the value of the publication for the scientific community. Our focus is specifically on two indicators, namely the impact factor of the journal in which a publication has appeared and the number of citations a publication has received during the first year after its appearance. Other indicators that could be used are the number of downloads of a publication (Brody, Harnad, & Carr, 2006), the number of readers according to a service such as Mendeley (Thelwall & Wilson, 2015), and other types of altmetric indicators (Costas, Zahedi, & Wouters, 2015). In this paper, however, our focus is on impact factors and early citations.

The use of early citations to predict long-term citation impact has been studied in various papers. Burrell (2003), Glänzel (1997), Mingers (2008), Mingers and Burrell (2006), and Wang et al. (2013) propose mathematical models that describe how publications accumulate citations over time. Using these models, they predict the citation impact of a publication in the longer term based on the publication's short-term citation history. Adams (2005), Bornmann et al. (2013), Levitt and Thelwall (2011), and Wang (2013) present empirical analyses of the correlation between short-term and long-term citation counts. Based on an analysis of publications from 1993 in six fields in the physical and life sciences, Adams (2005) concludes that "across reasonably large samples of research publications (not individual papers) it is possible to use initial citation counts predictively to index emerging quality relative to the field" (p. 579). Wang (2013) performs an analysis of all publications from 1980 indexed in the Web of Science database and reports that the Spearman correlation between short-term citation counts and citation counts after 31 years "rises from 0.266 in year 1 to 0.756 in year 3, and then slowly reaches 1 in year 31" (p. 866). Studies based on correlations reveal general patterns. Individual publications may of course strongly deviate from these patterns. Extreme deviations can be observed in the case of so-called 'sleeping beauties', which are publications that are hardly cited for a long time and then suddenly receive a lot of citations (Van Raan, 2004). The phenomenon of sleeping beauties illustrates the difficulty of making accurate predictions of long-term citation impact.

We are aware of three studies in which a comparison is made between the use of early citations and the use of impact factors for predicting longer-term citation impact. Abramo, D'Angelo, and Di Costa (2010) compare rankings of Italian universities based on citations and based on impact factors. They find that in certain fields, in particular in mathematics and in computer sciences, the ranking based on impact factors outperforms the ranking based on early citations in terms of the correlation with the ranking based on longer-term citation impact. Levitt and Thelwall (2011) propose a combined indicator of the impact of a publication that is obtained by taking a weighted average of the number of citations of a publication and the impact factor of the journal in which the publication has appeared. They report that in the case of a citation window of no more than one year the combined indicator provides a better prediction of the longer-term citation impact of publications in the field of economics than a straightforward indicator based only on citations. These results are in line with the findings of Stern (2014) for publications in the fields of economics and political science. Stern (2014) reports that shortly after the appearance of a publication the combined use of early citations and impact factors yields a better prediction of the longer-term citation impact of the publication than the use of early citations only.

We have now provided an overview of the literature that is most closely related to the research that we present in this paper. To make clear how our research contributes to the literature, let us summarize how our research differs from existing work:

- Our interest is in predicting long-term citation impact based exclusively on impact factors and early citations. As mentioned above, we do not want to use variables that can be easily manipulated or that may cause self-reinforcing effects.
- Our interest is in predicting long-term citation impact within one or two years after the appearance of a publication. Unlike some earlier studies (Levitt & Thelwall, 2008; Wang et al., 2011, 2012, 2013), we do not want to wait for five or more years before making predictions.
- Earlier work has shown that predicting long-term citation impact is a difficult task. Hence, it cannot be expected that the future number of citations of a publication can be predicted with a high degree of accuracy. Unlike most earlier work, our interest therefore is in predicting a probability distribution for the future number of citations of a publication. This probability distribution represents the uncertainty that we have about the number of citations a publication will receive in the future. We want this probability distribution to be predicted with a high degree of accuracy. We do not aim to provide a point estimate of the future number of citations of a publication.
- Most earlier studies (an exception is Glänzel, 1997) actually consider a simplified version of the problem of predicting long-term citation impact. For example, a study may consider the problem of predicting the number of citations that publications from 2005 have received by the end of 2014, where the prediction is based on information available at the end of 2006. However, how do we know whether the prediction model obtained for this problem will also work well when it is applied to a different time period? For example, will the model also work well to predict, based on information available at the end of 2014, the number of citations that publications from 2013 will have received by the end of 2022? This essential question is left unanswered in most earlier studies, but it will be addressed in our research.

3. Data

In this paper, we use the in-house version of the Web of Science database of the Centre for Science and Technology Studies of Leiden University. Only publications of the document types 'article' and 'review' are included in the analysis. For counting citations, author self-citations are excluded.

In order to estimate the coefficients of the regression model presented in Section 4, we use a specific set of publications. Because citation behavior differs between fields of science, only publications in the field of physics that were published in 1984 are included. In order to be included, a publication must belong to at least one of the following Web of Science subject categories: *Applied Physics, Fluids and Plasma Physics, Atomic, Molecular and Chemical Physics, Multidisciplinary Physics, Condensed Matter Physics, Nuclear Physics, Particles and Fields Physics*, and *Mathematical Physics*. Our entire data set includes 56,207 publications.

As already explained, we build a regression model with two predictors: the number of early citations of a publication and the impact factor of the journal in which a publication was published. In the rest of this paper, we will refer to these predictors as covariates. The number of early citations is defined as the number of citations received by a publication in the first year after its appearance. It is denoted by c_1 . In our data set, c_1 is the number of citations that a publication has received before the end of 1985. Hence, for counting early citations, all citations received by a publication in 1984 and 1985 are included. The impact factors that we use are impact factors for the year in which a publication appeared, which is 1984 in our data set. The impact factor of a journal in 1984 is calculated as the average number of citations that publications published in the journal in 1982 and 1983 received in 1984. In the calculation of the impact factor of a journal, only publications of the document types 'article' and 'review' are taken into account, both on the citing side and on the cited side.

In Section 5.6, we also consider two sets of publications in the field of physics published respectively in 1990 and 2000. These sets of publications are used to evaluate the predictive performance of our regression model. In Section 7, we address the sensitivity of the regression coefficients to a specific field. To this end, we use a set of publications in the fields of biology and chemistry published in 1984.

4. Regression model for quantile prediction

As pointed out in Sections 1 and 2, our interest is not in providing a point estimate of the future number of citations of a publication. Instead, our focus is on predicting a probability distribution for the future number of citations of a publication. More specifically, our aim is to predict the quantiles of this probability distribution.

Formally, the p th quantile $q(p)$ of a random variable Y with distribution function f is given by

$$q(p) = F^{-1}(p) = \inf\{y : F(y) \geq p\}.$$

Hence, saying that a publication scores at the p th quantile means that the number of citations of the publication is greater than or equal to the number of citations of a proportion of p of all publications.

Our goal is to predict quantiles for the distribution of the number of citations received by a publication starting from the second year after its publication date. For example, in our data set of publications published in 1984, we consider quantiles for the number of citations received by a publication between January 1986 and December 2013. We refer to this as the future number of citations of a publication or, alternatively, as the long-term citation impact of a publication. In this section, we propose a model that predicts the quantiles of the long-term citation distribution of a publication, conditioned on the impact factor and the number of citations in the first year.

4.1. Models for quantiles

Like Ke (2013), we assume that each publication has a fitness factor η . This fitness factor gives information about the competitiveness of a publication relative to other publications in obtaining citations. The fitness factor depends on different factors ϕ_s that contribute to the success of a publication. The fitness factor is assumed to be a product of each of these factors raised to some power δ_s :

$$\eta \propto \prod_s \phi_s^{\delta_s}. \quad (1)$$

We predict the quantiles of the distribution of future citations, conditioned on the fitness factor. A higher fitness factor means that a publication has a higher competitiveness to obtain citations. Therefore one expects that the quantiles of publications with a higher fitness factor are higher. We assume that the p th quantile of publications with fitness factor η , denoted by $q(p|\eta)$, is proportional to η :

$$q(p|\eta) = \tilde{C}_p \eta.$$

Here \tilde{C}_p is a constant independent of η for each quantile p .

We consider three definitions of the fitness factor: a definition based only on the impact factor IF , a definition based only on the number of citations in the first year c_1 , and a definition based on both IF and c_1 . For clarity of notation, the exponents

Table 1
Three models studied in this paper.

Model	η	Quantile prediction
Only IF	$\eta \propto IF^\beta$	$q(p IF) = \tilde{C}_p IF^{\beta p}$
Only c_1	$\eta \propto (c_1 + k_0)^\gamma$	$q(p c_1) = \tilde{C}_p (c_1 + k_0)^{\gamma p}$
Full model	$\eta \propto IF^\beta (c_1 + k_0)^\gamma$	$q(p IF, c_1) = \tilde{C}_p IF^{\beta p} (c_1 + k_0)^{\gamma p}$

δ_1 and δ_2 from Eq. (1) are relabeled as β and γ . The three definitions of the fitness factor are summarized in Table 1. The constant k_0 is needed to account for publications that have zero citations after one year. We will discuss our choice for k_0 in Section 5.1.

4.2. Quantile regression

In the model described in Section 4.1, the logarithm of the quantiles is linear in the logarithm of the covariates. For example, for the full model in Table 1 we obtain

$$\ln(q(p|IF, c_1)) = \gamma_p \ln(c_1 + k_0) + \beta_p \ln(IF) + C_p, \quad (2)$$

where $C_p = \ln(\tilde{C}_p)$. Because the logarithm is an increasing function, the logarithm of the p th quantile is equal to the p th quantile of the log-transformed citation counts. This means that we can take the logarithm of the number of citations and then fit Eq. (2) to the quantiles of those values.

Eq. (2) is fitted using quantile regression introduced by Koenker and Bassett (1978). While in standard least squares regression the sum of squared errors is minimized, quantile regression minimizes a different function. It solves

$$\min_{\xi} \sum_i \rho_p(y_i - x_i \xi). \quad (3)$$

Here y_i is the logarithm of the future number of citations of publication i , x_i is the vector of log-transformed covariates corresponding to publication i , $\xi^T = [C_p \ \beta_p \ \gamma_p]$, and the function ρ_p is defined as

$$\rho_p(z) = zp - z \mathbb{1}_{z < 0} = \begin{cases} zp & \text{if } z \geq 0 \\ z(p-1) & \text{if } z < 0. \end{cases}$$

Eq. (3) minimizes the difference between p and the fraction of negative residuals (Koenker & Bassett, 1978). Hence, when all values are different, the empirical quantiles for the future number of citations are fitted precisely. In our case, many publications have the same number of citations, and therefore there can be small differences between actual and fitted quantiles (for an illustration, see the blue dots in Fig. 11 below).

5. Quantile regression results

In this section, we apply the regression model from Section 4 to the data described in Section 3.

5.1. Model coefficients

From a research evaluation perspective, high quantiles of citation distributions are especially important because interest often focuses on identifying high-impact research. For this reason, in our analysis we consider the 0.5th up to the 0.99th quantile. Fig. 1a–c shows the parameters C_p , β_p and γ_p resulting from the quantile regression. The coefficients are shown for the three different versions of the model listed in Table 1.

The coefficient C_p in Fig. 1a is increasing in p for all three versions of the model. This is to be expected, because the quantiles are nondecreasing in p . We also see that C_p is a convex function. For the higher quantiles, C_p grows faster in p than for the lower quantiles. This indicates that for example the 0.98th and the 0.99th quantile are further away from each other than, say, the 0.60th and the 0.61th quantile. In Section 6, we will explain the behavior of C_p more precisely using quantile estimators for Pareto-tailed distributions.

The coefficients β_p and γ_p are decreasing in p . Hence, the impact factor and the number of citations in the first year have less influence on the long-term citation impact of highly cited publications than on the long-term citation impact of publications with an average number of citations. Another observation is that the parameter γ_p in the model with only c_1 follows a wave-like pattern: γ_p first increases as p increases, but at some point γ_p then decreases as p increases. For small values of p , a similar pattern can be observed for C_p in the model with only c_1 . This is caused by the fact that there are only a couple of different values of c_1 in the data set, since most publications have only a small number of citations in the first year. This makes the parameters that are fitted unstable in p .

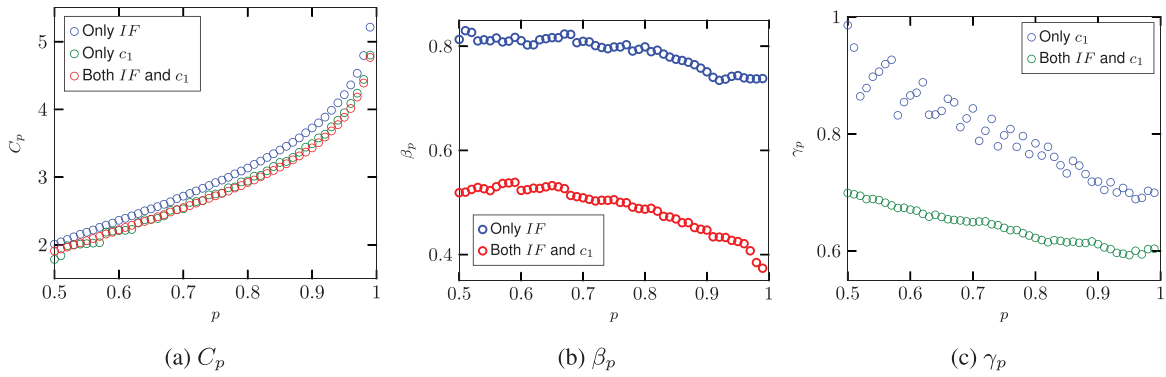


Fig. 1. Quantile regression coefficients for p th quantile versus p for the different models.

5.2. Influence of k_0

The parameter k_0 in the models listed in Table 1 is not fitted in the quantile regression. To get an understanding of the influence of k_0 on the regression coefficients, quantile regression is used to obtain the coefficients C_p , β_p , and γ_p for several values of k_0 . Fig. 2a–c shows the values of the regression coefficients for $k_0 = 0.3$ to 1.5 for the full model. We see that k_0 has hardly any influence on β_p . Also, k_0 does not have much influence on γ_p and C_p . Essentially, these coefficients increase or decrease by a constant value if k_0 changes. We use the value of k_0 that minimizes the sum of the squared difference between the fraction of publications with fewer citations than the predicted p th quantile and p . This results in $k_0 = 0.5$.

5.3. Fit of the models

We now investigate the fit of the models listed in Table 1 to the data. The fit of the model that uses only the impact factor is illustrated in Fig. 3. For every value of the impact factor, the predicted 0.50th, 0.80th, and 0.95th quantiles are shown as solid lines and the empirical quantiles are shown as dots. To create this figure, all impact factors are rounded to one decimal. A group of publications with the same rounded impact factor is included in the figure if there are at least 50 publications in the group. It is clear that the predicted and empirical quantiles may differ a lot. In a similar manner, the fit of the model that uses only the number of early citations is illustrated in Fig. 4. For the 0.50th quantile, the empirical and the predicted quantiles almost overlap. For the 0.80th and the 0.95th quantile, the model fits well for publications with a small number of citations in the first year, but it underestimates the quantiles for publications with a large number of early citations.

To illustrate the fit of the full model, we plot the predicted quantiles against the empirical quantiles. To do so, we first create groups of publications. A group consists of publications that all have the same c_1 and the same impact factor, where impact factors have been rounded to halves. Figs. 5–7 show all groups that include at least 50 publications. The figures relate to, respectively, the 0.50th, 0.80th, and 0.95th quantile. Each dot in the figures corresponds to a group of publications with the same IF and c_1 . The 45-degree lines are shown as a reference. In the case of a perfect fit, all dots should be located on the 45-degree lines. We see that for many groups of publications predictions are quite accurate, but there are also quite some groups for which there is a large difference between the predicted and the empirical quantile. Taking a closer look at the data, we see that, naturally, a better fit is obtained for larger groups. For example, the 0.95th quantile in a group with 50 observations is equal to the third largest observation, whereas in a group of 500 observations it is equal to the 26th

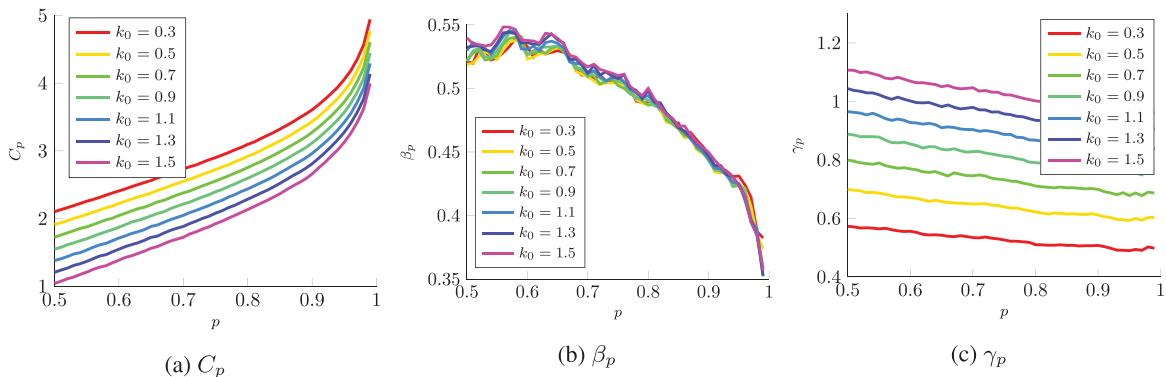


Fig. 2. Quantile regression coefficients for p th quantile versus p for k_0 ranging from 0.3 to 1.5.

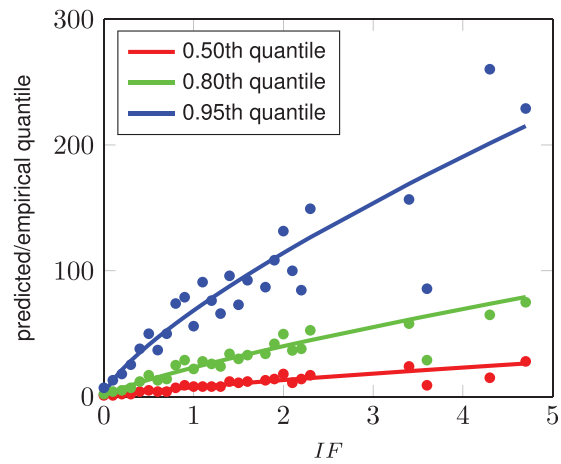


Fig. 3. Predicted value (solid line) and empirical value (dots) of the 0.50th, 0.80th, and 0.95th quantile versus IF for the model that uses only IF .

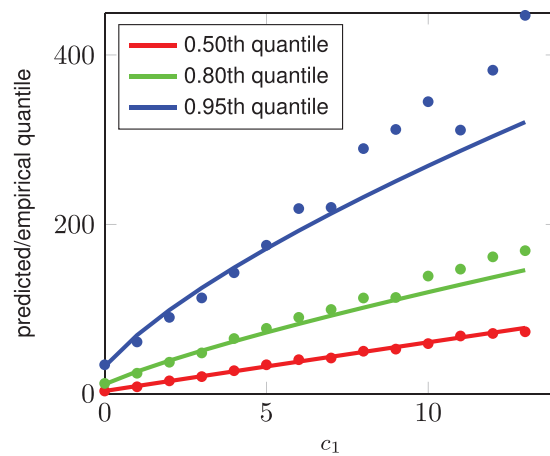


Fig. 4. Predicted value (solid line) and empirical value (dots) of the 0.50th, 0.80th, and 0.95th quantile versus c_1 for the model that uses only c_1 .

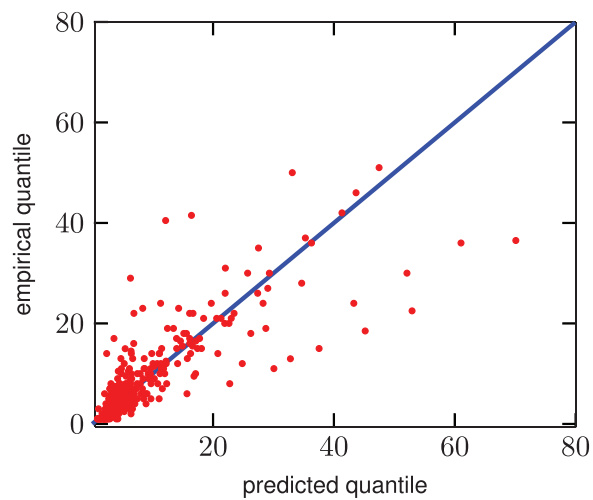


Fig. 5. Predicted versus empirical 0.50th quantile for groups with at least 50 publications.

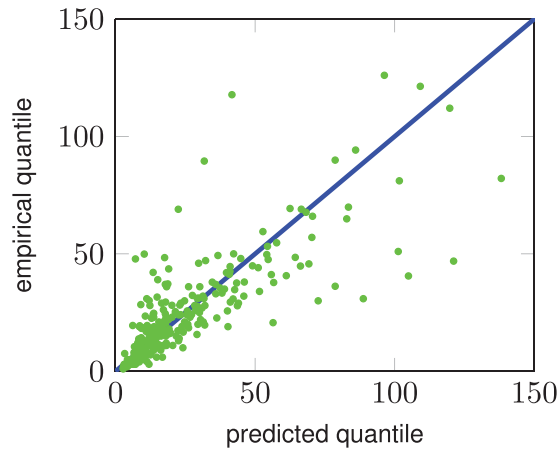


Fig. 6. Predicted versus empirical 0.80th quantile for groups with at least 50 publications.

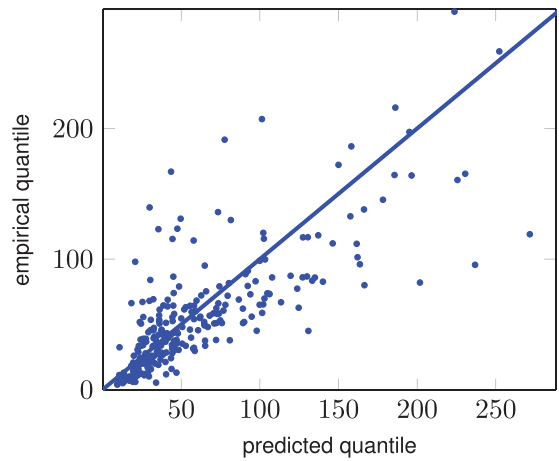


Fig. 7. Predicted versus empirical 0.95th quantile for groups with at least 50 publications.

largest observation. The third largest observation in a smaller group is more easily influenced by a few outliers than the 26th largest observation in a larger group of publications. In Fig. 8, the predicted quantiles are again plotted against the empirical quantiles, but only groups including at least 500 publications are shown. The 0.50th, 0.80th, and 0.95th quantiles are presented in the same plot. For the 0.50th and 0.80th quantile, the fit is excellent. For the 0.95th quantile, there are

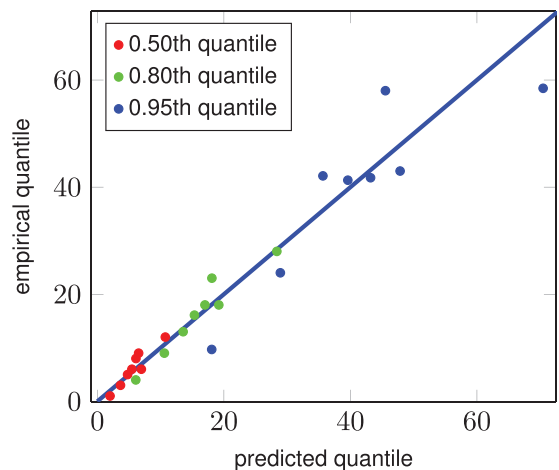


Fig. 8. Predicted versus empirical 0.50th, 0.80th, and 0.95th quantile for groups with at least 500 publications.

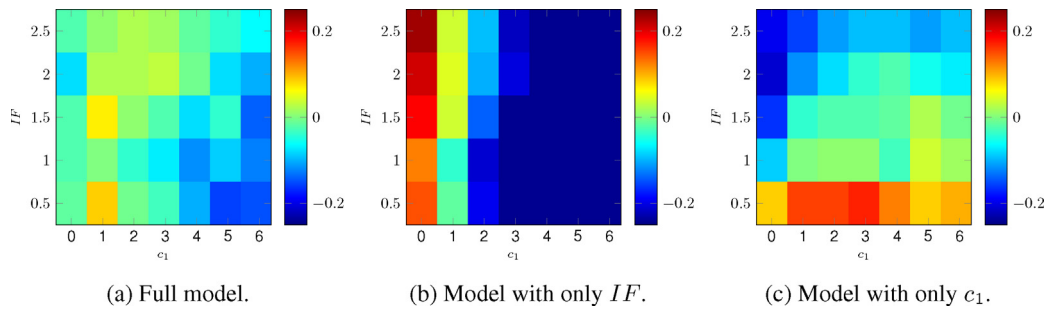


Fig. 9. Let f be the fraction of publications with fewer citations than the predicted 0.50th quantile. The figures show $f - 0.5$ for the three models for different c_1 and IF .

more outliers. We note that estimates for high quantiles will be explored further using quantile estimators for Pareto tails in Section 6.

5.4. Comparing the fit of the different models

Fig. 9a–c illustrates the difference in the fit of the three models. Let f be the fraction of publications with fewer citations than the predicted 0.50th quantile. For different groups of publications, the figures show $f - 0.5$. If the model predicts correctly, we expect this value to be close to zero. This is indicated by a green color in the figures. Red colors represent positive values, indicating that the model overestimates the 0.50th quantile for a group of publications. Blue colors correspond to negative values, which means underestimation of the 0.50th quantile. Each rectangle in the figures represents a group of publications with the same impact factor and the same number of citations in the first year. For example, for the model that uses only the impact factor, the model overestimates the 0.50th quantile for publications with an impact factor of 2.5 and with 0 citations in the first year. It underestimates the 0.50th quantile for publications with an impact factor of 2.5 and with 5 citations in the first year.

Based on these figures, we see that the model which uses only the impact factor does not predict very well for publications with either a small or a large number of citations in the first year. Likewise, the model that uses only the number of early citations does not predict very well for publications with either a low or a high impact factor. Similar figures can be created for other quantiles instead of the 0.50th quantile. From these figures we conclude that the full model yields more accurate predictions than the other two models. This means that both the impact factor and the number of citations in the first year provide important information for predictive purposes, and that impact factors and early citations should therefore be used together to obtain accurate predictions. For this reason, in the remainder of the results, the full model is used.

5.5. Predicting the conditional citation distribution

Using the quantile regression model, we can predict the entire conditional distribution of the number of citations. Fig. 10a and b shows the predicted and empirical conditional distribution for publications that have impact factor zero and zero citations in the first year and for publications that have impact factor one and one citation in the first year respectively. It is

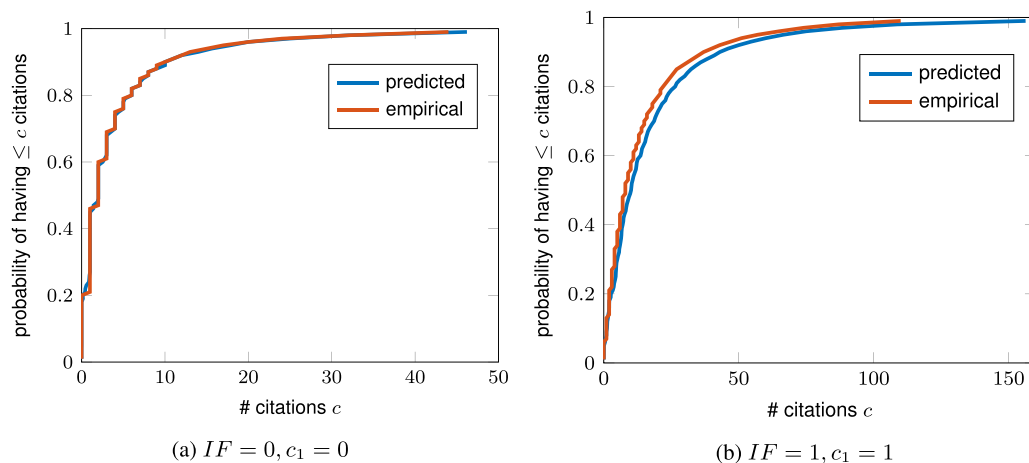


Fig. 10. Empirical and predicted conditional distribution function of the number of citations after 30 years.

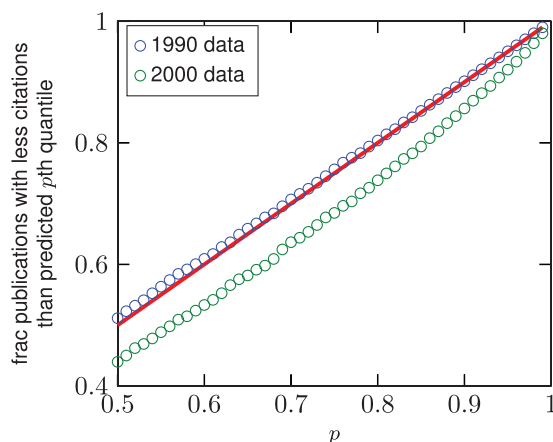


Fig. 11. p versus the fraction of publications with fewer citations than their predicted p th quantile. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

clear that publications with the same impact factor and the same number of citations in the first year, may have different citation numbers after 30 years. For this reason, predicting the entire conditional distribution is more valuable than giving a point estimate on the number of citations that publications receive. Furthermore, the conditional distributions in the two figures are different, which indicates that it is important to take into account the influence of the impact factor and the number of early citations. The quantile regression method predicts the conditional distribution quite accurately, especially for the publications with an impact factor of zero, and zero citations in the first year.

5.6. Predictions for later publications

In this section, we test whether the model fitted based on older publications also predicts well when applied to more recent publications. To this end, we first estimate the quantile regression coefficients for publications in the field of physics published in 1990. The model is fitted to predict the quantiles of the number of citations that these publications have received by the end of 2000. We then use the resulting model to predict the quantiles of the number of citations that publications in the field of physics published in 2000 have received by the end of 2010.

The predictive performance of the model is illustrated in Fig. 11. We compute the fraction of publications that have received fewer citations than their predicted p th quantile. If the model predicts well, this fraction should be p . In Fig. 11, p is plotted against the fraction of publications with fewer citations than their predicted p th quantile. Results are shown both for publications from 1990 (which were used to fit the model) and for publications from 2000 (which were not used in model fitting). The 45-degree line is included as a reference. In the case of a perfect fit, all dots should be located on the 45-degree line.

For publications from 1990, the quantiles are predicted almost perfectly, which is to be expected because we use quantile regression (see Section 4.2). However, when the model is applied to publications from 2000, we see that the quantiles are underestimated. For example, only around 43% of the publications from 2000 have received fewer citations than their predicted 0.50th quantile. This disappointing result must be due to structural changes that have taken place over time and that cause a model fitted to older data not to perform well when applied to newer data. In particular, there is a trend to include more and more references in publications, and as a result of this trend, the average number of citations that publications receive has increased over time (Wallace, Larivière, & Gingras, 2009). Because the model is fitted based on older publications, which have lower citation counts than more recent publications, the model underestimates the quantiles for more recent publications.

We want to adjust the predictions of the model for the increase over time in the average number of citations per publication. To do so, we make predictions based on normalized data. This means that all inputs and outputs of the model are divided by their average value. For example, the number of citations of a publication in the first year, c_1 , is divided by the average value of c_1 over all publications. Similarly, the number of citations of a publication after 10 years is divided by the average number of citations after 10 years over all publications. The quantile regression model is fitted on the normalized data from 1990. The data from 2000 are normalized in the same way, and the model fitted based on publications from 1990 is used to predict the quantiles for publications from 2000. The resulting predictions are normalized predictions with respect to the average number of citations after 10 years over all publications from 2000. Hence, a predicted quantile of for example 2 means that the quantile equals twice the average number of citations after 10 years over all publications from 2000.

Fig. 12 shows that the performance of the model for publications from 2000 has become just as good as for publications from 1990. Like in Fig. 11, p is plotted against the fraction of publications with fewer citations than their predicted p th quantile. The good performance of the model based on normalized data indicates that the increasing number of citations

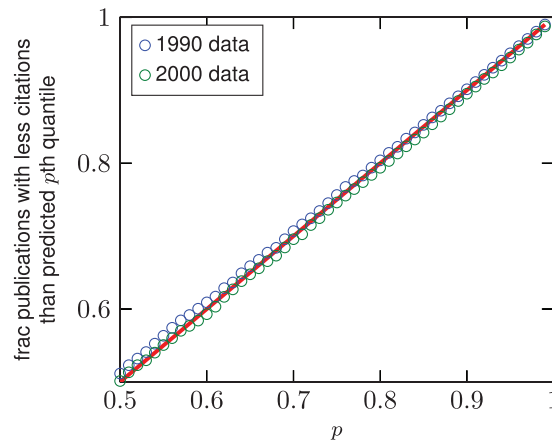


Fig. 12. p versus the fraction of publications with fewer citations than their predicted p th quantile. An adjustment has been made for increasing citation counts over time.

received by publications is indeed responsible for the underestimation that can be observed in Fig. 11. The use of normalized data solves this problem.

6. Tail quantiles

In this section, we take a closer look at high quantiles using a quantile estimation technique from extreme value theory.

6.1. Tail of the citation distribution

In the literature, a lot of attention has been paid to the tail of citation distributions, and in particular to the question whether this tail follows a Pareto distribution. The possibility of a Pareto tail was already suggested by De Solla Price (1976). Redner (1998) analyzed a large data set of publications and their citations and observed that the tail of the distribution of citations over publications can be described by a power law. Clauset, Shalizi, and Newman (2009) proposed a statistical methodology for testing the presence of power-law behavior in empirical data. Based on the same data set as Redner (1998), they concluded that for citation distributions a power-law tail cannot be ruled out. The methodology of Clauset et al. (2009) was also used by Albarrán, Crespo, Ortuño, and Ruiz-Castillo (2011), who found that in a large number of scientific fields citation distributions seem to have a power-law tail.

Formally, let X be the number of future citations. The random variable X has a Pareto tail if for some x_1 we have

$$P(X > x) = wx^{-\alpha}, \quad x \geq x_1.$$

Here w and α are parameters. α is also called the tail exponent.

There are many ways to test whether a distribution has a Pareto tail. Here we use the Zenga plot proposed by Cirillo (2013). The motivation behind this method is that it allows distinguishing between a Pareto tail and a lognormal tail, while many other methods, such as the QQ-plot, fail to detect this difference.

Let f be the distribution function of a random variable X . The Zenga curve Z is defined as

$$Z(u) = 1 - \frac{Q^-(u)}{Q^+(u)}, \quad 0 < u < 1,$$

$$Q^-(u) = \frac{1}{u} \int_0^u F^{-1}(s) ds, \quad 0 \leq u \leq 1,$$

$$Q^+(u) = \frac{1}{1-u} \int_u^1 F^{-1}(s) ds, \quad 0 \leq u \leq 1.$$

Hence, the Zenga curve is a measure of how much weight of the distribution lies below the u th quantile relative to how much weight lies above the u th quantile, as a function of u . The Zenga curve has different shapes for different distributions. For the lognormal distribution, the Zenga curve is a straight line, while for Pareto distributions, the Zenga curve is a convex increasing function (Cirillo, 2013).

Figs. 13 and 14 show the empirical Zenga plot for groups of publications with either the same number of citations in the first year or the same impact factor. The Zenga curves are clearly convex increasing functions. This indicates that for all groups of publications the citation distribution has a Pareto tail.

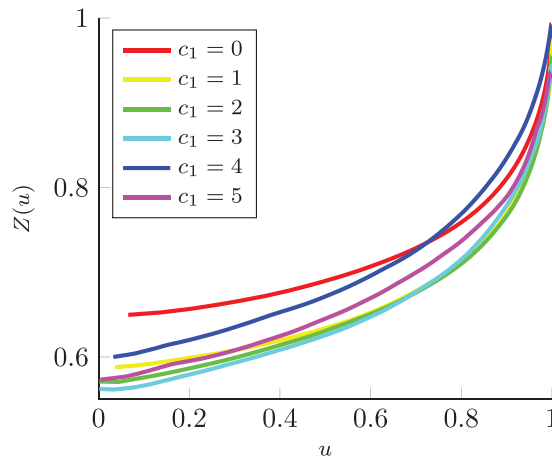


Fig. 13. Zenga plot for different c_1 .

6.2. Pareto quantile estimator

Under the assumption of a Pareto tail, we can estimate high quantiles using the estimator proposed by Dekkers et al. (1989). For a group of publications with given IF and c_1 , this estimator is of the form

$$q(p|IF, c_1) = X_{(n-k,n)} \left(\frac{k}{n(1-p)} \right)^{1/\alpha} \tag{4}$$

Here $n := n(IF, c_1)$ is the number of publications in the group of publications with given IF , and c_1 and $X_{(n-k,n)} := X_{(n,n-k)}(IF, c_1)$ is the number of citations of the k th most cited publication in this group. Furthermore, $k := k(IF, c_1)$ is the threshold where the Pareto tail starts. So the k publications with the largest number of citations follow a Pareto distribution.

The Pareto tail starts at the $(1 - (k/n))$ th quantile. This threshold quantile will be called the p^* th quantile. Thus, $X_{(n-k,n)}$ is the empirical value of the p^* th quantile. The tail index α is estimated using the Hill estimator (Hill, 1975). The threshold k and hence the threshold quantile p^* is estimated using the procedure suggested by Beirlant, Glänzel, Carbonez, and Leemans (2007). This procedure minimizes an approximation of the asymptotic mean squared error of the estimate of α . The resulting estimate for p^* is $p^* = 0.95$.

Fig. 15 plots the quantiles that are predicted by Eq. (4) against the empirical quantiles, when $X_{(n-k,n)}$ is given. Note that the accuracy of the predictions is inconclusive because we rely on $X_{(n-k,n)}$, which is not known in practice. However, we see that Eq. (4) captures the behavior of high quantiles quite accurately. In the next section, we will obtain insight into the predictions for high quantiles by linking the Pareto estimator to the quantile regression estimator.

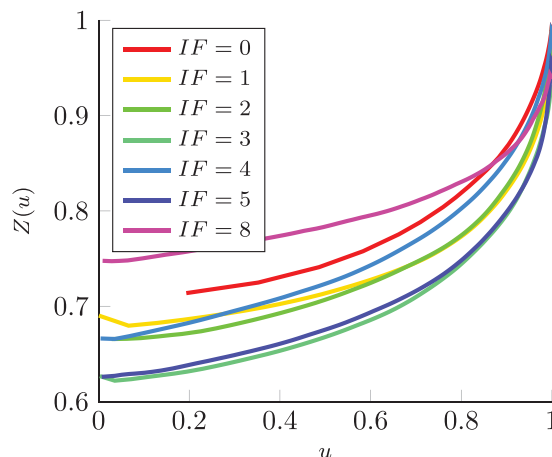


Fig. 14. Zenga plot for different IF .

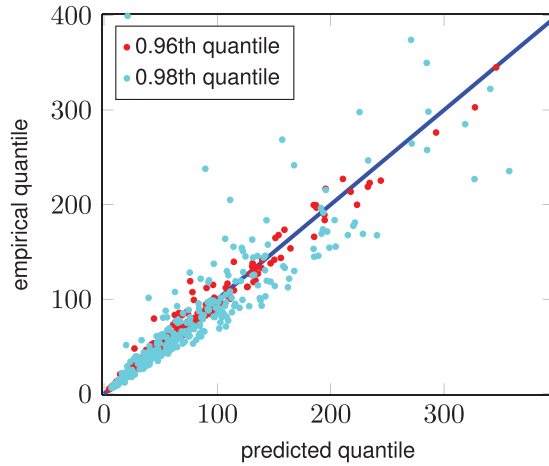


Fig. 15. Predicted p th quantile obtained from the Pareto estimator versus empirical p th quantile for $p=0.96$ and $p=0.98$ and for groups with at least 50 publications.

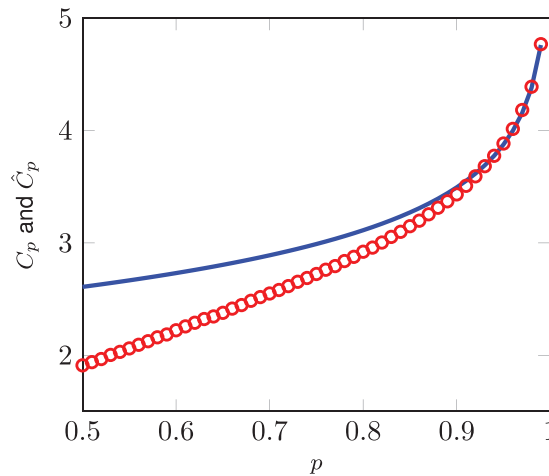


Fig. 16. \hat{C}_p from Eq. (6) (line) and C_p from regression (dots) for p th quantile.

6.3. Linking the Pareto and regression estimators

As mentioned above, the problem with Eq. (4) is that it uses $X_{(n-k,n)}$, the empirical value of the p^* th quantile, at which the Pareto tail starts. In practice, $X_{(n-k,n)}$ is not known and needs to be predicted. A natural way to overcome this is to replace $X_{(n-k,n)}$ by the predicted p^* th quantile from the quantile regression:

$$q(p^*|IF, c_1) = \tilde{C}_{p^*} IF^{\beta_{p^*}} (c_1 + k_0)^{\gamma_{p^*}}.$$

We then obtain the following estimator for the tail quantiles:

$$q(p|IF, c_1) = \left(\frac{1-p^*}{1-p}\right)^{1/\alpha} \tilde{C}_{p^*} IF^{\beta_{p^*}} (c_1 + k_0)^{\gamma_{p^*}}, \quad p \geq p^*, \tag{5}$$

where we used the identity $(k/n) = 1 - p^*$.

We can now explain the behavior of $C_p = \ln(\hat{C}_p)$ in Fig. 1a by comparing Eq. (5) to the regression estimator. If the two estimators were equal, then for $p \geq p^*$ we would have

$$\left(\frac{1-p^*}{1-p}\right)^{1/\alpha} \tilde{C}_{p^*} IF^{\beta_{p^*}} (c_1 + k_0)^{\gamma_{p^*}} = \tilde{C}_p IF^{\beta_p} (c_1 + k_0)^{\gamma_p}.$$

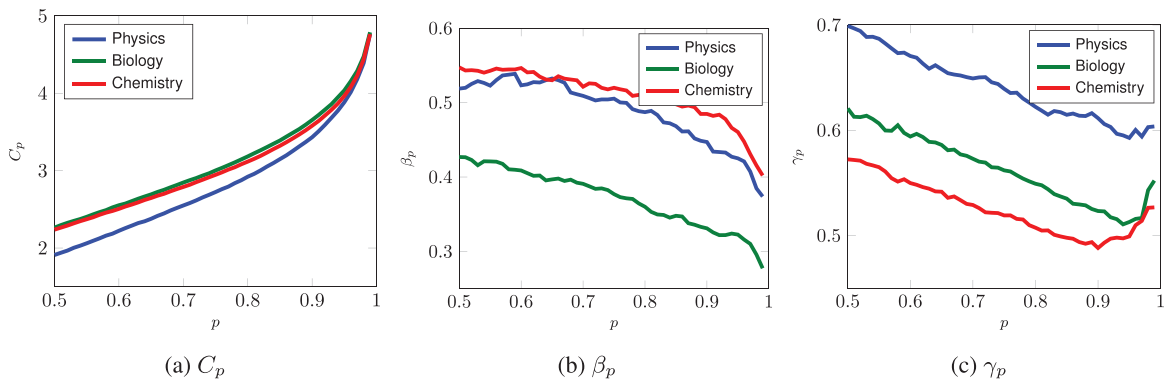


Fig. 17. Quantile regression coefficients for p th quantile versus p for three fields of science.

Hence, if we assume that β_p and γ_p are constant for large values of p , this suggests that we may use $((1-p^*)/(1-p))^{1/\alpha} \tilde{C}_{p^*}$ for \tilde{C}_p . By taking the logarithm, we obtain the following proxy \hat{C}_p for the regression coefficient C_p :

$$\hat{C}_p = C_{p^*} + \frac{1}{\alpha} (\ln(1-p^*) - \ln(1-p)), \quad p \geq p^*. \quad (6)$$

In Fig. 16, C_p and \hat{C}_p are plotted. The line corresponds to \hat{C}_p given by Eq. (6) with $p^* = 0.95$. For completeness, we plot this line for all $p \in [0.5, 0.99]$. The dots correspond to the regression coefficients C_p . We see that Eq. (6) indeed can be used as an analytical description of C_p when $p \geq p^*$. For $p \geq 0.95$, there is an excellent agreement between C_p and \hat{C}_p . In fact, Eq. (6) shows a good agreement for all $p \geq 0.9$.

7. Sensitivity of parameters to the field of science

In the previous sections, data from the field of physics was used to fit the coefficients of the regression model. In this section, we study the influence of the field of science on the regression coefficients. We again consider publications published in 1984, and the quantile regression is again used to predict the conditional quantiles of the distribution of the number of citations received by these publications by the end of 2013. The coefficients C_p , β_p , and γ_p obtained from the quantile regression are plotted in Fig. 17a–c for three different fields: Biology, chemistry, and physics.

The regression coefficient C_p is higher for publications in biology and chemistry than for publications in physics. Other things being equal, this means that publications in biology and chemistry receive more citations than publications in physics. This is the case mainly for the lower quantiles. For the higher quantiles, the differences in the coefficient C_p are small. The regression coefficient β_p is lower for biology publications than for publications in physics and chemistry. This means that in biology the impact factor is less determining for the long-term citation impact of publications. However, the differences in this coefficient are small. The coefficient γ_p is highest for publications in physics. Hence, in physics the number of citations that a publication has received in the first year is more determining for the publication's long-term citation impact than in biology and chemistry.

8. Conclusions

We have proposed a model to predict a probability distribution for the future number of citations of a publication. Two predictors are considered in the model: The impact factor of the journal in which a publication has appeared and the number of citations received by a publication in the first year after its appearance. The proposed model is based on quantile regression. The good fit of the model indicates that quantile regression is a suitable tool to predict the quantiles of the probability distribution of a publication's future number of citations. We have found that the quantile regression coefficients β_p and γ_p , corresponding to respectively the impact factor and the number of early citations, are not stable in the quantile p . Hence, the influence of the impact factor and the number of early citations on the long-term citation impact of a publication is different for different quantiles.

Three variants of our prediction model have been studied. The variant in which both the impact factor and the number of early citations are used turns out to fit the data better than the variants in which only one of the two predictors is included. This means that both the impact factor and the number of early citations are important to predict the probability distribution of a publication's future number of citations.

Importantly, our proposed model provides accurate predictions also for publications that were published later than the publications used for estimating the model parameters. However, in order to obtain these accurate predictions, it is necessary to normalize all inputs and outputs of the model by their average value.

We emphasize that we obtain accurate predictions of the probability distribution for the future number of citations of a publication, not of the future number of citations itself. This is a key difference. We are able to accurately predict the probability distribution for the future number of citations of a publication, but this probability distribution may turn out to be very wide, in which case we still face a lot of uncertainty about the number of citations a publication will receive in the future.

We have also investigated the tail of the citation distributions obtained in our analysis. Zenga plots (Cirillo, 2013) have been used for this purpose. It turns out that the tail of our citation distributions can be approximated by a Pareto distribution. Using an estimator for the tail quantiles of a Pareto distribution (Dekkers et al., 1989), we have obtained an explicit equation for the regression coefficient C_p in our model for high quantiles p .

There are a number of issues that require further research. First of all, further research may focus on the fitness factor that we use in our model. Following Ke (2013), we have assumed that the fitness factor is a product of our two predictors each raised to a certain power. Other ways of modeling the fitness factor may also be investigated.

The analysis presented in this paper is based on publications in the field of physics. This is a broad field consisting of many different subfields. These subfields probably all have their own citation practices. Differences in citation practices between fields or subfields have not been taken into account in our prediction approach. Further research may focus on linking our prediction approach to the literature on field normalization of citation-based indicators.

Another issue for further research is the use of other predictors, in addition to impact factor and early citations. In Section 2, we already suggested some possibilities: Number of downloads of a publication, number of readers according to a service such as Mendeley, and other types of altmetric indicators. Further research may investigate the effect of adding these predictors to our model. In particular, it would be interesting to find out whether the use of additional predictors decreases the level of uncertainty in predictions of long-term citation impact.

Finally, perhaps the most challenging issue for further research is to make predictions of long-term citation impact not only for individual publications but also for the entire publication oeuvre of a researcher, a research group, or a research institution (e.g. Acuna, Allesina, & Kording, 2012; Bornmann, 2013). Moving from predictions at the individual publication level to predictions at the level of oeuvres of publications is an essential step in the development of a methodology for bibliometric research evaluation based on very recent publications, but this step is far from trivial. A prediction approach that yields accurate results at the individual publication level may provide biased results when it is used at the level of oeuvres of publications.

References

- Abramo, G., D'Angelo, C., & Di Costa, F. (2010). Citations versus journal impact factor as proxy of quality: Could the latter ever be preferable? *Scientometrics*, 84(3), 821–833.
- Acuna, D., Allesina, S., & Kording, K. (2012). Future impact: Predicting scientific success. *Nature*, 489(7415), 201–202.
- Adams, J. (2005). Early citation counts correlate with accumulated impact. *Scientometrics*, 63(3), 567–581.
- 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(2), 385–397.
- Beirlant, J., Glänzel, W., Carbonez, A., & Leemans, H. (2007). Scoring research output using statistical quantile plotting. *Journal of Informetrics*, 1(3), 185–192.
- Bornmann, L. (2013). The problem of citation impact assessments for recent publication years in institutional evaluations. *Journal of Informetrics*, 7(3), 722–729.
- Bornmann, L., Leydesdorff, L., & Wang, J. (2013). Which percentile-based approach should be preferred for calculating normalized citation impact values? an empirical comparison of five approaches including a newly developed citation-rank approach (p100). *Journal of Informetrics*, 7(4), 933–944.
- Brody, T., Harnad, S., & Carr, L. (2006). Earlier web usage statistics as predictors of later citation impact. *Journal of the American Society for Information Science and Technology*, 57(8), 1060–1072.
- Burrell, Q. (2003). Predicting future citation behavior. *Journal of the American Society for Information Science and Technology*, 54(5), 372–378.
- Cirillo, P. (2013). Are your data really Pareto distributed? *Physica A*, 392(23), 5947–5962.
- Clauet, A., Shalizi, C., & Newman, M. (2009). Power-law distributions in empirical data. *SIAM Review*, 51(4), 661–703.
- Costas, R., Zahedi, Z., & Wouters, P. (2015). Do “altmetrics” correlate with citations? Extensive comparison of altmetric indicators with citations from a multidisciplinary perspective. *Journal of the Association for Information Science and Technology* (in press).
- De Solla Price, D. J. (1976). A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science*, 27(5), 292–306.
- Dekkers, A. L., Einmahl, J. H., & De Haan, L. (1989). A moment estimator for the index of an extreme-value distribution. *Annals of Statistics*, 17(4), 1833–1855.
- 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.
- Fu, L., & Aliferis, C. (2010). Using content-based and bibliometric features for machine learning models to predict citation counts in the biomedical literature. *Scientometrics*, 85(1), 257–270.
- Glänzel, W. (1997). On the possibility and reliability of predictions based on stochastic citation processes. *Scientometrics*, 40(3), 481–492.
- Haslam, N., Ban, L., Kaufmann, L., Loughnan, S., Peters, K., Whelan, J., et al. (2008). What makes an article influential? Predicting impact in social and personality psychology. *Scientometrics*, 76(1), 169–185.
- Haslam, N., & Koval, P. (2010). Predicting long-term citation impact of articles in social and personality psychology. *Psychological Reports*, 106(3), 891–900.
- Hill, B. M. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics*, 3(5), 1163–1174.
- Ke, W. (2013). A fitness model for scholarly impact analysis. *Scientometrics*, 94(3), 981–998.
- Koenker, R., & Bassett, G., Jr. (1978). Regression quantiles. *Econometrica*, 46(1), 33–50.
- Levitt, J., & Thelwall, M. (2008). Patterns of annual citation of highly cited articles and the prediction of their citation ranking: A comparison across subjects. *Scientometrics*, 77(1), 41–60.
- Levitt, J., & Thelwall, M. (2011). A combined bibliometric indicator to predict article impact. *Information Processing and Management*, 47(2), 300–308.
- Lokker, C., McKibbin, K., McKinlay, R., Wilczynski, N., & Haynes, R. (2008). Prediction of citation counts for clinical articles at two years using data available within three weeks of publication: Retrospective cohort study. *BMJ*, 336(7645), 655–657.

- Mingers, J. (2008). Exploring the dynamics of journal citations: Modelling with s-curves. *Journal of the Operational Research Society*, 59(8), 1013–1025.
- Mingers, J., & Burrell, Q. (2006). Modeling citation behavior in management science journals. *Information Processing and Management*, 42(6), 1451–1464.
- Mingers, J., & Xu, F. (2010). The drivers of citations in management science journals. *European Journal of Operational Research*, 205(2), 422–430.
- Onodera, N., & Yoshikane, F. (2015). Factors affecting citation rates of research articles. *Journal of the Association for Information Science and Technology*, 66(4), 739–764.
- Peters, H., & Van Raan, A. (1994). On determinants of citation scores: A case study in chemical engineering. *Journal of the American Society for Information Science*, 45(1), 39–49.
- Redner, S. (1998). How popular is your paper? An empirical study of the citation distribution. *European Physical Journal B*, 4(2), 131–134.
- Stern, D. (2014). High-ranked social science journal articles can be identified from early citation information. *PLOS ONE*, 9(11), e112520.
- Thelwall, M., & Wilson, P. (2015). Mendeley readership altmetrics for medical articles: An analysis of 45 fields. *Journal of the Association for Information Science and Technology*. (in press).
- Van Raan, A. (2004). Sleeping beauties in science. *Scientometrics*, 59(3), 467–472.
- Wallace, M. L., Larivière, V., & Gingras, Y. (2009). Modeling a century of citation distributions. *Journal of Informetrics*, 3(4), 296–303.
- Walters, G. (2006). Predicting subsequent citations to articles published in twelve crime-psychology journals: Author impact versus journal impact. *Scientometrics*, 69(3), 499–510.
- Wang, D., Song, C., & Barabási, A.-L. (2013). Quantifying long-term scientific impact. *Science*, 342(6154), 127–132.
- Wang, J. (2013). Citation time window choice for research impact evaluation. *Scientometrics*, 94(3), 851–872.
- Wang, M., Yu, G., Xu, J., He, H., Yu, D., & An, S. (2012). Development a case-based classifier for predicting highly cited papers. *Journal of Informetrics*, 6(4), 586–599.
- Wang, M., Yu, G., & Yu, D. (2011). Mining typical features for highly cited papers. *Scientometrics*, 87(3), 695–706.
- Yu, T., Yu, G., Li, P.-Y., & Wang, L. (2014). Citation impact prediction for scientific papers using stepwise regression analysis. *Scientometrics*, 101(2), 1233–1252.