



Axiomatic measures of intellectual influence[☆]

Ignacio Palacios-Huerta^{b,*}, Oscar Volij^a

^a Ben-Gurion University of the Negev, Israel

^b London School of Economics and Ikerbasque Foundation at UPV/EHU, United Kingdom



ARTICLE INFO

Available online 12 March 2014

PsycINFO classification:

D70

D71

D89

Keywords:

Intellectual influence

Rankings

Axiomatic approach

ABSTRACT

Measuring influence allows the study of such issues as the impact and reputation of scientists and scientific publications, the dynamics of innovation, and the construction of ranking algorithms for search engines in the world wide web. Ranking methods that measure influence are typically based on the information contained in the network of communications between different entities (scholarly publications, patents, web pages). This paper presents within the same framework recent developments that use the axiomatic approach to derive ranking methods. Two related but essentially different ranking problems are studied: journal ranking problems and scholar ranking problems. The paper concludes with open theoretical and empirical questions for future research.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

In the past few years there has been a significant growth in the use of influence measures and numerical indices that quantify “impact” using the network of citation data between entities such as scholarly publications, patents, web pages, and other forms of documentation.

The information contained in these networks is valuable because it allows us to make a first attempt at a rigorous quantitative analysis of elusive phenomena. In academic and scientific environments, for instance, these data are useful to measure reputation, the quality of journals, and the productivity of scholars, universities and others. Citation analysis is in fact widely used as a management tool for making decisions on hiring, promotion, salary, etc., in academic institutions. In the literature on the dynamics of innovation and technological change, patent citation data are used for examining the pattern of knowledge spillovers and evaluating the importance of private and public patenting (see Jaffe and Trajtenberg (2002)). And in the world wide web, citation

analysis plays a fundamental role in the development of search algorithms that find and rank web pages.

Beginning at least with Garfield's (1972) impact factor, there was a proliferation of ranking methods in economics and other social sciences. Unfortunately this early literature made little effort to justify the use of one method over another. Rather, the justification was simply based on intuitive grounds or on the impression that they yielded introspectively plausible results. Posner (2000, p. 383) lamented these shortcomings by arguing that “citation analysis is not an inherently economic methodology: most of it has been conducted without any use of the theories or characteristic techniques of economists”.

This situation has somewhat changed in recent years when several studies began to apply the axiomatic methodology to address the problem of measuring intellectual influence. Instead of proposing ranking methods on intuitive grounds, the idea is to derive them from first principles. Thus, this methodology allows us to characterize and compare different measures according to the properties that they satisfy.

The literature on intellectual influence has developed two related but essentially different branches of problems: journal ranking problems and scholar ranking problems. This paper combines them into a single framework and presents some of the more interesting axiomatic results which have been derived over the last few years. Needless to say, this presentation is not exhaustive. We conclude with a few open questions.

2. Framework for ranking problems

In what follows we use the terms “article” and “journal” to refer to the cited unit of publication and to the citing unit, respectively. We also use the word “citations” to refer to the citations obtained by an

[☆] We thank Tommaso Valletti, Marco Ottaviani and participants in seminars at the universities of Zurich and East Anglia, and at the 40th Annual Conference of EARIE in Évora for useful comments. Financial support from the Ministerio de Economía y Competitividad (ECO2012-31626, ECO2012-31346), and from the Departamento de Educación, Política Lingüística y Cultura of the Gobierno Vasco (IT869-13, IT568-13) is gratefully acknowledged.

* Corresponding author at: Department of Management, London School of Economics, Houghton Street, London WC2A 2AE, London, UK.

E-mail addresses: i.palacios-huerta@lse.ac.uk (I. Palacios-Huerta), ovolij@bgu.ac.il (O. Volij).

article, and “references” to mean citations made by a journal. As indicated earlier, other interpretations can be given to the problem of measuring influence based on communication data depending on the specific application.

Within the academic context, a typical CV, as will likely be required by tomorrow’s university, may look like this:

	JPE	AER	JET	IJIO
Art a	2	3	0	1
Art b	5	2	4	3
Art c	1	7	0	2

Here, each row represents a published article and each column represents a journal. An entry represents the number of citations that the corresponding article obtains from a paper published in the corresponding journal. Thus, for instance, Article *b* has been cited three times by the papers published in *IJIO* while *AER* made seven references to Article *c*.

One can identify in the literature two approaches to evaluating a CV like this one. According to the first approach, which is typically applied when the CV belongs to a journal, all articles are considered equal and the references of each journal are condensed into a single number, namely the total number of references made by the journal. Thus, the relevant data are exhibited as

JPE	AER	JET	IJIO
8	12	4	6

According to the second approach, which is typically applied when the CV belongs to an individual scholar, the citations obtained by a given article are considered equal and are usually condensed into a single number, namely its total number of citations. Thus, the relevant data are exhibited as

Art a	6
Art b	14
Art c	10

These two approaches lead to two different ranking problems, journal ranking problems and scholar ranking problems, which are analyzed next.

3. Journal ranking problems

A journal ranking problem consists of a group of journals and their respective citation records. For instance, in this journal ranking problem:

	IJIO	AER	JPE
IJIO	100	100	150
AER	50	200	100
JPE	100	50	150

the first row represents the *IJIO*’s citation record and the first column represents the *IJIO*’s “opinion” about the various journals.

Formally, a journal ranking problem consists of a set of journals $\mathcal{J} = \{1, \dots, J\}$ and a $J \times J$ irreducible matrix C . The entry c_{ij} is the number of citations that Journal i received from Journal j . Each column represents the corresponding journal’s *opinion* about the journals in \mathcal{J} . The *reference intensity* of Journal $j \in \mathcal{J}$ is the number of citations it issues, namely $r_j = \sum_{i \in \mathcal{J}} c_{ij}$. In any given journal ranking problem, different journals may have different reference intensities.

The objective is to take a journal ranking problem and aggregate the journals’ opinions into an “objective opinion”. Formally, we are interested in a ranking method f that takes a journal ranking problem (\mathcal{J}, C) and returns a list of scores $v = (v_j)_{j \in \mathcal{J}}$, which represents the objective opinion. As with prices, the absolute values v_j are meaningless. Only

the relative values v_i/v_j are important. For this reason, if one needs a unique list of values, some normalization is required. In what follows, if two lists of scores $v = (v_i)_{i \in \mathcal{J}}$ and $w = (w_i)_{i \in \mathcal{J}}$ are proportional to each other, we will write $v \propto w$.

3.1. Examples of ranking methods

In what follows, we describe a number of ranking methods.

Impact Factor Method. This method is given by the function that assigns to each ranking problem (\mathcal{J}, C) an opinion proportional to the total sum of citations it obtained:

$$\left(\sum_{j \in \mathcal{J}} c_{1j}, \dots, \sum_{j \in \mathcal{J}} c_{ij} \right).$$

Invariant Method. This method is given by the function that assigns to each ranking problem $C = \{c_{ij}\}_{(ij) \in \mathcal{J}^2}$ an opinion $(v_i)_{i \in \mathcal{J}}$ that satisfies

$$v_i = \sum_{j \in \mathcal{J}} \frac{c_{ij}}{r_j} v_j \quad i \in \mathcal{J}.$$

That is, the Invariant Method assigns to each journal a weighted sum of its citations, where the weights are the scores of the citing journals divided by their respective reference intensities.

Handicap Method. This method is given by the function that assigns to each ranking problem $C = \{c_{ij}\}_{(ij) \in \mathcal{J}^2}$ an opinion $(v_i)_{i \in \mathcal{J}}$ that satisfies

$$v_i = \sum_{j \in \mathcal{J}} \frac{c_{ij}}{q_j} \quad i \in \mathcal{J}$$

$$q_j = \sum_{i \in \mathcal{J}} \frac{c_{ij}}{v_i} \quad j \in \mathcal{J}$$

The Handicap Method issues for each journal i two values. One is its score $-v_i$, and the other is its “competence” $-1/q_i$. The score measures the impact of the journal on other journals and depends on how much it is cited. Its “incompetence” $-q_i$, measures the influence he gets from other journals, and depends on how much it cites. The score of the journal according to this method is the sum of its citations weighted by the citing journal’s competence. And the incompetence of a journal is the weighed sum of its references where the weights are the inverse of the cited journal’s score.

3.2. Special problems

It is sometimes convenient to focus attention not on all the general problems, possibly arising from complicated networks of citations, but only on some simple ones. The following are two such problems.

Normalized problems. These are problems in which all the journals have the same reference intensity. That is, for some r

$$r_j = \sum_{i \in \mathcal{J}} c_{ij} = r \quad j \in \mathcal{J}.$$

Row-balanced problems. These are problems in which all journals have the same number of citations. That is, for some c

$$\sum_{j \in \mathcal{J}} c_{ij} = c \quad i \in \mathcal{J}.$$

3.3. Axioms

We now list various axioms against which different ranking methods can be compared. The first two axioms are based on the simple problems just described.

Uniformity. A ranking method is *uniform* if it assigns the same score to all the journals whenever the ranking problem is both normalized and row-balanced.

Exactness. A ranking method is *exact* if whenever it assigns the same score to all the journals and the ranking problem is normalized, the problem is also row-balanced.

Uniform methods assign equal scores to all the journals if the problem is both normalized and row-balanced. Exact methods, on the other hand, assign equal scores to all the journals in a normalized problem only if it also row-balanced. It can be checked that all the methods introduced in Section 3.1 satisfy these two axioms.

The next axiom requires that the reference intensities do not affect the rankings.

Invariance to reference intensity. Let $P = \langle \mathcal{J}, C \rangle$ be a ranking problem and let $P' = \langle \mathcal{J}, C' \rangle$ be the normalized problem that is obtained from P by dividing each journal's opinion by its reference intensity. That is, $c'_{ij} = c_{ij}/r_j$ for all $i, j \in \mathcal{J}$. A ranking method f is *invariant to reference intensity* if for all such two problems the same opinion is produced, namely, if $f(P) \propto f(P')$.

This axiom states that a journal cannot affect the ranking by simply citing more. In order to do so it must change its opinion about the journals, that is the relative distribution of its references. It is easy to check that the Invariant and Handicap methods satisfy this axiom, whereas the Impact Factor method does not.

The next axiom, which we call *homogeneity* requires that if every journal multiplies its references to journal i^* by some common factor λ , then the score of journal i^* must also be multiplied by λ . That is, if every journal increases its opinion about journal i^* by a factor of λ , then the relative score of this journal should also go up by λ .

Homogeneity. Let $P = \langle \mathcal{J}, C \rangle$ be a ranking problem and let $f(P) = (v_1, \dots, v_j)$. Let $i^* \in \mathcal{J}$ be a journal. Let $P' = \langle \mathcal{J}, C' \rangle$ be the ranking problem that is obtained from P by multiplying all the references to i^* by some constant factor $\lambda > 0$. That is, i^* 's row becomes $(c'_{i^*j}) = (\lambda c_{i^*j})$. We say that the ranking method f satisfies *homogeneity* if $f(P') \propto (v_1, \dots, v_{i^*-1}, \lambda v_{i^*}, v_{i^*+1}, \dots, v_j)$.

The next axiom, which is based on Stigler, Stigler and Friedland (1995), requires that in normalized two-journal problems, the relative scores must be the ratio of mutual citations.

Import–export principle. The ranking method f satisfies the *import–export principle* if for every normalized two-journal problem

$$C = \begin{pmatrix} c-a & b \\ a & c-b \end{pmatrix}$$

we have $f(C) \propto (b, a)$.

Before we present the final axiom, we need to define the concept of a *reduced problem*. Consider the ranking problem

$$C = \begin{pmatrix} 30 & 16 & 10 \\ 17 & 35 & 20 \\ 15 & 9 & 30 \end{pmatrix}$$

and suppose that this 3 by 3 matrix C is too large for our computer to work with but that, on the other hand, our computer is able to compute the scores for any two-journal ranking problem. If we were interested in the scores of Journals 1 and 2, the question is how we could build a 2 by 2 matrix that faithfully represented the network C . Clearly, deleting the third column and the third row

would mean ignoring relevant information. The reduced problem with respect to journal 3 takes into account the information that is contained in the third column and third row, and incorporates it into the first two columns and rows respecting journal 3's opinions. Specifically, the reduced problem with respect to Journal 3 is

$$C_3 = \begin{pmatrix} 30+5 & 16+3 \\ 17+10 & 35+6 \end{pmatrix}.$$

As we can see, in order to build C_3 , the 15 references of Journal 1 to Journal 3 have been redistributed back to Journals 1 and 2 in proportion to the opinions that Journal 3 has about them. Similarly, the 9 references of Journal 2 to Journal 3 are redistributed back to Journals 1 and 2 proportionally to the opinions of Journal 3 about them.

More generally, given a ranking problem $P = \langle \mathcal{J}, C \rangle$ its *reduced problem with respect to journal $i^* \in \mathcal{J}$* is given by $P' = \langle \mathcal{J} \setminus \{i^*\}, C_{i^*} \rangle$ where for all $i, j \in \mathcal{J} \setminus \{i^*\}$

$$c'_{ij} = c_{ij} + c_{i^*i} \frac{c_{i^*j}}{\sum_{j \neq i^*} c_{i^*j}}.$$

The *consistency* axiom requires that the relative scores of the Journals i and j ($\neq i^*$) be the same in both P and P' .

Consistency. Let $P = \langle \mathcal{J}, C \rangle$ be a ranking problem, let $i^* \in \mathcal{J}$ be a journal, and let $P' = \langle \mathcal{J} \setminus \{i^*\}, C_{i^*} \rangle$ be the reduced problem with respect to i^* . Let $f(P) = (v)_{i \in \mathcal{J}}$ and $f(P') = (v')_{i \in \mathcal{J} \setminus \{i^*\}}$. A ranking method f satisfies *consistency* if for all such ranking problems $(v)_{i \in \mathcal{J} \setminus \{i^*\}} \propto (v')_{i \in \mathcal{J} \setminus \{i^*\}}$.

The axioms we have presented seem reasonable requirements for ranking methods to satisfy. For a more detailed description of their desirability see Palacios-Huerta and Volij (2004) and Demange (2014). We can now present the main characterization results of this section:

Theorem 1. Demange (2014) *The handicap method is the only one that satisfies invariance to reference intensity, uniformity and homogeneity.*

Theorem 2. Demange (2014) *The handicap method is the only one that satisfies invariance to reference intensity, exactness and homogeneity.*

Theorem 3. (Palacios-Huerta and Volij (2004)) *The invariant method is the only one that satisfies invariance to reference intensity, the import–export principle, and consistency.*

It is not difficult to show that the Handicap method does not satisfy the import–export principle and consistency, whereas the Invariant method does not satisfy homogeneity. Interestingly, despite these differences, the illustrations in Demange (2014) show that these two methods generate somewhat similar rankings, especially for the top economic journals.

4. Scholar ranking problems

A *scholar* is a list $S = (c_1, \dots, c_n)$ of non-negative numbers, ordered from highest to lowest, where n is the number of publications and for each $i = 1, \dots, n$, c_i is an index of the citations obtained by publication i .

The following are three examples of scholars with five, six and three publications respectively:

$$S_1 = (6, 6, 5, 1, 0)$$

$$S_2 = (5, 3, 2, 2, 2, 1)$$

$$S_3 = (4, 4/3, \sqrt{2})$$

The index of citations, as in the case of scholar S_3 , need not be a natural number. It ought to be an index that measures the impact of the corresponding publication. It can be, for instance, its citations




Fig. 1. Two scholars.

weighted according to any of the methods described in the previous section. The first two scholars are depicted graphically in Fig. 1.

As with journal ranking problems, the objective is to order different scholars according to some objective scoring function. However, while in journal ranking problems the relative scores matter, in scholar ranking problems only the order of the scholars matters.

Given two scholars, $S = (c_1, \dots, c_{n_1})$ and $S' = (c'_1, \dots, c'_{n_2})$, we can define the *best scholar* among the (weakly) worse scholars than both of them as follows:

$$S \wedge S' = (\min(c_1, c'_1), \dots, \min(c_{\min(n_1, n_2)}, c'_{\min(n_1, n_2)}))$$

Fig. 2 depicts $S_1 \wedge S_2$.

Similarly, given two scholars, $S = (c_1, \dots, c_{n_1})$ and $S' = (c'_1, \dots, c'_{n_2})$, we can define the *worst scholar* among the (weakly) better scholars than both of them as follows:

$$S \vee S' = (\max(c_1, c'_1), \dots, \max(c_{\max(n_1, n_2)}, c'_{\max(n_1, n_2)}))$$

Fig. 3 depicts $S_1 \vee S_2$.




Fig. 2. $S_1 \wedge S_2$.

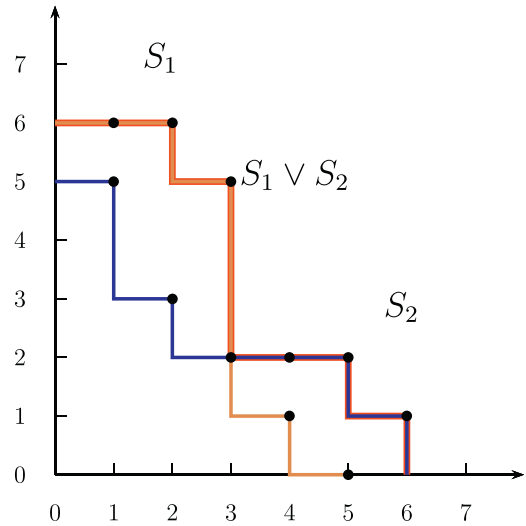


Fig. 3. $S_1 \vee S_2$.




Fig. 4. Two step-based indices.

Using this framework we next study two types of indices: step-based indices and measure-based indices.

4.1. Step-based indices

Chambers and Miller (2014) introduced the class of step-based indices. Any non-decreasing line from the origin defines a step-based index as follows: the index of a scholar S is the length of the line that is “inside” S ’s frontier. Fig. 4 shows two such lines and we can see that the index defined by the left line ranks S_1 above S_2 and the index defined by the right line ranks S_2 above S_1 .

Some well-known indices are step-based.¹ For example, the Hirsch index, $H(S_i)$, (Hirsch (2005)) is the step-based index defined by the 45 degree ray from the origin.² The Publication Count $N(S_i)$, which is simply the number of publications, is defined by the line that coincides with the horizontal axis. The Maximum index, $M(S_i)$, which is given by the article with most citations, is defined by the line that coincides

¹ Most indices in practice use raw citations as the scholars’ entries.
² The Hirsch index has been axiomatically characterized by Woeginger (2008) and Marchant (2009).




Fig. 5. Examples of step-based indices.

with the vertical axis. Similar to the i10 index (number of publications with at least 10 citations) provided by Google Scholar, the $i3$ index, $i3(S_i)$, is defined by the line that coincides with the vertical axis until the point (0,3) and from then on with the horizontal line with ordinate 3. Fig. 5 depicts these indices.

Chambers and Miller (2014) propose the following axiom:

Consistency with best and worst scholars. A scholar ranking index I satisfies consistency with best and worst scholars if for any two scholars S_1 and S_2 ,

$$I(S_1) \geq I(S_2) \Leftrightarrow \begin{cases} I(S_1) = I(S_1 \vee S_2) \\ I(S_2) = I(S_1 \wedge S_2) \end{cases}$$

This axiom requires that given two scholars S_1 and S_2 such that S_1 is ranked at least as high as S_2 , the best out of the (weakly) worse scholars than both of them must be ranked as high as the latter, and the worst out of the (weakly) better scholars than both of them must be ranked as low as the former.

Consistency with Best and Worst Scholars fully characterizes the class of step-based indices:

Theorem 4. Chambers and Miller (2014) A scholar ranking index satisfies Consistency with best and worst scholars if and only if it is a monotonic transformation of a step-based index.

4.2. Measure-based indices

Any measure on the positive orthant defines a measure-based index as follows: the index of a scholar S is the measure of the space under S 's frontier.

The citation count is an example of a measure-based index. Indeed, if we use the Lebesgue measure, the area under the graph of the scholar is exactly the total number of citations.

Another example is the product of two measures: one measure defined on the horizontal axis and another measure on the vertical axis. That is, take two increasing functions: μ , defined on the horizontal (articles) axis, and σ , defined on the vertical (citations) axis, such that the measure of any rectangle $[a,b] \times [c,d]$ is $(\mu(b) - \mu(a))(\sigma(d) - \sigma(c))$. For instance, letting μ be the identity function and $\sigma(x) = x^\alpha$ for some $\alpha > 0$, the “area” under the graph of any scholar $S = (c_1, \dots, c_n)$ is

given by

$$I(S) = \sum_{i=1}^n (c_i)^\alpha.$$

Perry and Reny (2013) propose the following axioms:

Dummy paper. $(c_1, \dots, c_n) \sim (c_1, \dots, c_n, 0)$

This axiom dictates that adding a paper with no citations to a scholar does not affect the ranking of the scholar.

Monotonicity. $(c_1, \dots, c_n) \geq (c'_1, \dots, c'_n) \Rightarrow (c_1, \dots, c_n) \succcurlyeq (c'_1, \dots, c'_n)$

Monotonicity requires that adding citations to the papers of a given scholar cannot decrease the scholar's ranking.

Independence. $(c_1, \dots, c_n) \succcurlyeq (c'_1, \dots, c'_n) \Rightarrow (c_1, \dots, c_n) \parallel (c) \succcurlyeq (c'_1, \dots, c'_n) \parallel (c)$

Here, \parallel denotes the concatenation operation with the further reordering in non-increasing order of the resulting list. Independence requires that adding one paper with the same number of citations to two scholars does not alter the scholars' relative rankings.

Scale invariance. For all $\lambda > 0$, $(c_1, \dots, c_n) \succcurlyeq (c'_1, \dots, c'_n) \Rightarrow (\lambda c_1, \dots, \lambda c_n) \succcurlyeq (\lambda c'_1, \dots, \lambda c'_n)$.

Scale invariance says that the relative ranking of two scholars is not affected if all the citations are multiplied by a given positive constant.

Quality responsiveness. For any $c_1 > 0$, $(c_1 + 1) > (c_1, 1)$

Finally, quality responsiveness holds that, for the case of scholars with a single publication, adding a citation to an already cited paper increases the index more than adding a publication with a single citation. Note that assuming independence, the requirement of quality responsiveness is automatically extended to all scholars, regardless of their number of publications. Further assuming Dummy paper, this axiom requires that a citation affects the index more when added to an already cited paper than when added to a publication with no citations.

Perry and Reny (2013) use these axioms to prove the following characterization:

Theorem 5. Perry and Reny (2013) A citation index satisfies Monotonicity, Dummy paper, Independence, Scale invariance and Quality responsiveness, if and only if it is a monotonic transformation of

$$I(S) = \sum_{i=1}^n (c_i)^p$$

for some $p > 1$.

This result says that the only indices that satisfy the above five axioms are positive monotonic transformations of the \mathcal{L}_p norm of the scholar, for some $p > 1$.

We conclude by mentioning several issues that are worth investigating. An open theoretical question is to find the axioms that may characterize not just the $I(S)$ index but the complete family of measure-based indices. An open empirical question is to study how well these different indices may align with labor market outcomes in the different academic markets (see Ellison (2013) for a study using the H -index in Economics). The same idea can be pursued for other markets where measuring influence using a network of communication data may be important. For instance, it can be applied to understand the determinants of public and private patenting or the demand for rankings (search engines) in the world wide web.

References

Chambers, C.P., Miller, A.D., 2014. Scholarly influence. J. Econ. Theory (forthcoming).
Demange, G., 2014. A ranking method based on handicaps. Theor. Econ. (forthcoming).

- Ellison, G., 2013. How does the market use citation data? The Hirsch index in economics. *Am. Econ. J.: Appl. Econ.* 5 (3), 63–90.
- Garfield, E., 1972. Citation analysis as a tool in journal evaluation. *Science* 178, 471–479.
- Hirsch, J.E., 2005. An index to quantify an individual's scientific research output. *Proc. Natl. Acad. Sci.* 102, 16569–16572.
- Jaffe, A.B., Trajtenberg, M., 2002. *Patents, citations and innovations*. MIT Press, Cambridge, MA.
- Marchant, T., 2009. An axiomatic characterization of the ranking based on the h-index and some other bibliometric rankings of authors. *Scientometrics* 80 (2), 325–342.
- Palacios-Huerta, I., Volij, O., 2004. The measurement of intellectual influence. *Econometrica* 72 (3), 963–977.
- Perry, M., Reny, P., 2013. How to count citations if you must. Mimeo, University of Warwick and University of Chicago.
- Posner, R.A., 2000. An economic analysis of the use of citations in the law. *Am. Law Econ. Rev.* 2 (2) (381–06).
- Stigler, G.J., Stigler, S.M., Friedland, C., 1995. The journals of economics. *J. Polit. Econ.* 103, 331–359.
- Woeginger, G., 2008. An axiomatic characterization of the Hirsch-index. *Math. Soc. Sci.* 56, 224–232.