Scoring of Web Pages and Tournaments - Axiomatizations*

Giora Slutzki† and Oscar Volij‡

April 11, 2004

Abstract

Consider a set of elements which we want to rate using information about their bilateral relationships. For instance sports teams and the outcomes of their games, journals and their mutual citations, web sites and their link structure, or social alternatives and the tournament derived from the voters' preferences. A wide variety of scoring methods have been proposed to deal with this problem. In this paper we axiomatically characterize two of these scoring methods, variants of which are used to rank web pages by their relevance to a query, and academic journals according to their impact. These methods are based on the Perron-Frobenius theorem for non-negative matrices.

Keywords: Perron-Frobenius theorem, scoring, rankings, tournaments, web pages, citations.

*We thank Herbert A. David for his insightful comments and an anonymous referee for several useful references.

†Department of Computer Science, Iowa State University, Ames, IA 50011, USA.
Email: slutzki@cs.iastate.edu.

‡Department of Economics, Iowa State University, Ames, IA 50011, USA.
Email: oscar@volij.co.il. http://volij.co.il.

1 Introduction

There are many instances in which several alternatives need to be assigned scores. A typical example is a round-robin tournament, where several contenders play against each other. (See, for instance, Moon (1968) and Laslier (1997)). Chess players are ranked according to the ELO scoring system. (See Elo (1978)). Citation indexes like the Journal of Citation Reports database, are used for ranking scientific journals by their intellectual impact. (See Pinski and Narin (1976) and Palacios-Huerta and Volij (2002)). In social choice theory, social alternatives are ranked based on voters' preferences. (See Arrow (1963), Sen (1970) and Moulin (1988) who emphasize impossibility results, and Saari (2001) and Laslier (1997) where different methods for ranking candidates in elections are studied.)

In the last decade, with the advent of the World Wide Web (WWW), numerous web-search engines developed efficient techniques for mining information from the internet in response to user queries. Because of the huge amounts of data on the WWW, it is of paramount importance that data mined by the search engines be well organized and submitted back to the user in some order of relevance to the query asked. The newer search engines view the WWW as a directed graph which is analyzed relative to the given query using methods based on the Perron-Frobenius theory of the eigenvectors of non-negative matrices. (See, for instance, Brin and Page (1998), Kleinberg (1999), Page, Brin, Motwani, and Winograd (1999), and Chakrabarti, Dom, Gibson, Kleinberg, Kumar, Raghavan, Rajagopalan, and Tomkins (1999).)[1]

All of these problems share the same structure: they consist of a set of "players" (alternatives, journals, web pages, etc.) and a data matrix that summarizes the relationships among the players. In a generalized tournament, the entries of the matrix are the number of victories of one player over the other; in a journal ranking problem the entries may be the number of citations that each journal gets in the others; in a web page ranking problem, the entries represent the number of links from one page to the other; in a social choice problem, the entries represent the number of voters that prefer one alternative over the other.

Much work has been done on scoring systems and ranking methods. See for example

[1]An anonymous referee suggested Saaty's work on the analytic hierarchy process in the context of multicriteria decision making as another successful application of the Perron-Frobenius theory.

Zermelo (1929), Wei (1952), Kendall (1955), Daniels (1969), Moon and Pullman (1970), Kano and Sakamoto (1985), David (1987), Keener (1993), Levchenkov (1995), van den Brink and Gilles (2000), Herings, van der Laan, and Talman (2001), Palacios-Huerta and Volij (2002), Conner and Grant (2000), and the comprehensive texts of David (1988) and Laslier (1997). The simplest scoring method that is used regularly for ranking participants in sports tournaments is the *points* method, which assigns to each participant the number of victories he obtained. This ranking system is often referred to as the Copeland score and it was axiomatized by Rubinstein (1980), Henriet (1985), and van den Brink and Gilles (2000).[2]

In this paper we axiomatically characterize two scoring methods that were proposed by Daniels (1969) and Moon and Pullman (1970), and which are based on the Perron-Frobenius theorem for non-negative matrices. These methods, which we call the *fair-bets* and the *Invariant* scoring methods, have a nice intuitive description based on the idea of fairness in betting systems. A list of numbers $v = (v_1, ..., v_n)$ may be viewed as a system of bets in the following sense. A bet on i is a contract whereby the bettor gets $\$v_j$ each time player i beats player j and he pays out $\$v_i$ each time player i loses regardless of i's opponent's identity. The bet on i is *fair* if the total payment due to i's losses equals the total revenue obtained from i's victories. The fair-bets method assigns scores $v = (v_1, \ldots, v_n)$ to the n players so that the resulting system of bets is fair. The Invariant scores, on the other hand, are (proportional to) the amounts that these fair bets pay as a result of each player's victories (which, by definition, are equal to the total amount paid due to each player's losses). In terms of web page scoring, the fair-bets score (v_i) represents the value of any reference from page i, whereas the invariant score of web page i is the sum of its citations from the other web pages, weighted by their respective fair-bets values.

As seen above, scoring problems can represent a wide variety of situations where some kind of ranking or scoring is needed. Correspondingly, a given scoring method may be more appropriate in some applications and not adequate in others. For instance, we will argue that the fair-bets scoring method is appropriate for scoring participants in tournaments (of all kinds), but is not adequate for ranking journals or web pages based on citations (hypertext links). On the other hand, we shall argue that the Invariant method

[2]Readers interested in learning some of the subtle properties of this procedure can consult Merlin and Saari (1996).

is appropriate for ranking journals and web pages, but is not adequate for tournaments.

The reader may wonder at this point, on what basis one scoring method can be said to be more appropriate or adequate than another. The best framework to deal with these kind of questions is by using the axiomatic approach. Namely, by shifting the discussion from the methods themselves, to the more fundamental level of the properties satisfied by them. We will say that a scoring method is adequate for some application if it satisfies properties reasonably required from scoring methods intended for that application.

The axioms we use for our first characterization of the fair-bets scoring method are uniformity, inverse proportionality to losses, and neutrality. Uniformity requires that if all players share the same number of wins and losses, they should be uniformly ranked. Inverse proportionality to losses requires that if in a particular scoring problem the players are equally ranked and one player's number of losses against each of his opponents is multiplied by a given constant, then his score, relative to the score of each one of his opponents, should be divided by the same constant. Neutrality dictates that if all players share the same number of wins and losses, and if the scoring function ranks them all equally, then it should still rank them all equally if any two teams play several more games against each other which they evenly split. We show that the only scoring function that satisfies these three properties is the fair-bets scoring function.

The Invariant function is characterized by means of uniformity, invariance to reference intensity and weak additivity. Invariance to reference intensity requires that the scores be invariant to multiplication of a web page's references by a constant. Weak additivity imposes that if for some very simple problems, those for which all web pages share the same number of citations and references, the function assigns scores that are proportional to the number of citations, then it should continue to score in proportion to the number of citations if a pair of web pages add the same number of mutual links.

Lastly, we provide an additional axiomatic characterization of the fair-bets scoring function for a class of problems with a variable number of players. The key properties are reciprocity and consistency. The former says that in two-player problems, the scores should be proportional to the players' respective number of victories. The latter is a property that requires some "consistency" between the ranking of a problem and its reduced problems.

The paper is organized as follows. After setting up the notation in Section 2, we present the model in Section 3.1. After introducing the fair-bets and the Invariant scoring functions, we discuss their intuitive interpretations in Section 3.2. Section 3.3 introduces several properties of scoring functions, some of which are used in Section 4.1 to characterize the fair-bets and the Invariant scoring functions. The independence of the axioms is shown in Section 4.2, and Section 4.3 provides the alternative characterization of the fair-bets scoring function based on the consistency principle.

2 Notation and Preliminaries

The set of natural numbers is denoted by $I\!N$ and $\mathcal{N}$ denotes the set of finite, nonempty subsets of $I\!N$. For each $N \in \mathcal{N}$ define $\Delta_N = \{(v_i)_{i\in N} : v_i \geq 0, \quad \sum_{i\in N} v_i = 1\}$. Further, let $\Delta = \bigcup_{N\in\mathcal{N}} \Delta_N$. For any vector $(\lambda_i)_{i\in N}$, $\text{diag}(\lambda_i)_{i\in N}$ denotes the diagonal matrix with $(\lambda_i)_{i\in N}$ as its main diagonal. We say that two vectors u and v are proportional, denoted $u \propto v$, if there is a positive real κ such that $u = \kappa v$. For any vector $v \in \mathbb{R}^n$, we denote $\|v\| = \sum_i |v_i|$. For any matrix $A = (a_{ij})$, we write a_{i*} for the row sum $\sum_j a_{ij}$, and a_{*j} for the column sum $\sum_i a_{ij}$. Also, we denote $C_A = \text{diag}(a_{*j})_{j\in N}$ the diagonal matrix whose diagonal entries are the column sums of A. We say that $j \in N$ *is reachable from* $i \in N$, if there is a finite sequence $i_0, \ldots, i_n$, with $i_0 = i$ and $i_n = j$ such that $\prod_{k=0}^{n-1} a_{i_k,i_{k+1}} > 0$. We say that i and j *communicate* if either $i = j$ or if i is reachable from j and j is reachable from i. It can be checked that the communication relation is an equivalence relation. Therefore it partitions N into equivalence classes. The matrix A is *irreducible* if N is the only element of this partition.

3 The model

3.1 Scoring problems and functions

Let $A = (a_{ij})_{(i,j)\in N\times N}$ be a finite nonnegative matrix. As noted in the previous section, the communication relation partitions the set N into equivalence classes. If N_1 and N_2 are two distinct equivalence classes, then one of the following mutually exclusive statements holds:

1. Every element of N_1 is reachable from every element of N_2 and no element of N_2 is reachable from any element of N_1.

2. Every element of N_2 is reachable from every element of N_1 and no element of N_1 is reachable from any element of N_2

3. No element of N_2 is reachable from any element of N_1 and no element of N_1 is reachable from any element of N_2.

Each of these three possibilities yields a natural ranking between elements of N_1 and elements of N_2: either one is reachable from the other, or they are incomparable. A non trivial problem is, however, the ranking and scoring of elements within equivalence classes, since there the elements are reachable from each other. This leads to the following definition.

A *scoring problem* is a pair $\langle N, A \rangle$, where $N \in \mathcal{N}$ is a finite set and A is an $|N| \times |N|$ irreducible non-negative matrix.

The set N contains the elements to which we want to assign scores. They can represent teams, social alternatives, web pages, journals, etc. depending on the interpretation of the scoring problems. In this paper, generic elements of N will be called players. The matrix A contains information about the relevant relations among the elements of N, based on which the scores will be assigned. Examples of scoring problems are tournaments, generalized tournaments, web page ranking problems, etc. In a generalized tournament, the matrix A is such that for all $i, j \in N$, $i \neq j$, $a_{ii} = 0$ and $a_{ij} + a_{ji} = m$ for some fixed m. In a tournament, which is a generalized tournament with $m = 1$, the matrix A is restricted to be a (0,1)-matrix. The entry a_{ij} in the generalized tournament matrix A is interpreted as the number of victories of player i over player j. In a web page ranking problem the members of N are interpreted as web pages and the entry a_{ij} of the "citation matrix" A is interpreted as the number links from web page j to web page i.[3] In order to avoid confusion between incoming and outgoing hypertext links, links to page i will be referred to as *i's citations*, and links from page i will be referred to as *i's references*. Sometimes we will use the term "link" to refer to either a citation or a reference. The entries of the matrix in a journal ranking problem have an analogous interpretation. One can also interpret a scoring problem more generally as a "tournament," like the problem of ranking a set of chess players, in which the players do not necessarily play the same number of times

[3]Note that in this context, A is the transpose of the adjacency matrix of the WWW graph.

against each other. Another interpretation would be the problem of ranking a set of social alternatives where the entry a_{ij} represents the number of voters that prefer alternative i to alternative j.

There are some problems that one may not want to consider. For example, in tournaments, participants do not play against themselves. Similarly, web pages seldom link to themselves, and if they do, one may not want to take these self-references into account. Also, in a social choice problem where voters have reflexive preferences, all voters weakly prefer any alternative to itself. For this reason, in this paper we will restrict attention to the sublcass $\mathcal{S}_0$ of scoring problems $\langle N, A\rangle$ such that for all $i \in N$, its main diagonal entry $a_{ii} = 0$.[4] In fact, in the first two results, we will restrict attention to the smaller subclass $\mathcal{S}_0(N) \subseteq \mathcal{S}_0$ of problems with a fixed player set N. All of our results will remain valid (with a slight modification in Section 4.3) for the larger class of all scoring problems. The interpretation of the axioms will not be as natural, though.

We are interested in associating numerical scores, $v_i \geq 0$, to all players in a scoring problem, reflecting the strength of the players relative to each other. Since only relative scores matter, we restrict attention to scoring vectors whose components add up to one.

A *scoring function* is a function that assigns to each scoring problem $S = \langle N, A\rangle \in \mathcal{S}_0$, a vector of scores $v \in \Delta_N$. Here are two examples.

The *Invariant* scoring function, I, is the function that assigns to each scoring problem $S = \langle N, A\rangle$, the unique scores $I(S) = v = (v_i)_{i\in N} \in \Delta_N$ that satisfy

$$v_i = \sum_{j\in N} \frac{a_{ij}}{a_{*j}} v_j \qquad \text{for all } i \in N, \tag{1}$$

or, in matrix notation,

$$v = AC_A^{-1}v. \tag{2}$$

The *fair-bets* scoring function, F, is the function that assigns to each scoring problem $S = \langle N, A\rangle$, the unique scores $F(S) = v = (v_i)_{i\in N} \in \Delta_N$ that satisfy

$$\sum_{j\in N} a_{ji} v_i = \sum_{j\in N} a_{ij} v_j \qquad \text{for all } i \in N, \tag{3}$$

[4]Note, however, that this restriction should not be applied for ranking journals, since self-references do matter. See Palacios-Huerta and Volij (2002).

or, in matrix notation,

$$C_A v = Av. \tag{4}$$

Note that for all scoring problems $S = \langle N, A\rangle$, $I(S) \propto C_A F(S)$. Indeed, v satisfies (4) if and only if $C_A v$ satisfies (2). These two scoring functions were introduced in the literature by Daniels (1969) and Moon and Pullman (1970). They are well-defined: the Invariant function selects the unique[5] stationary distribution of the irreducible stochastic matrix AC^{-1} and, as stated above, the fair-bets vector of scores $F(S)$ is proportional to $C_A^{-1} I(S)$. Since in every scoring problem $S = \langle N, A\rangle$ the matrix A is irreducible, both $I(S)$ and $F(S)$ are strictly positive vectors. The Invariant method is at the core of the method used by Google to rank web pages, known as PageRank (see for example Page, Brin, Motwani, and Winograd (1999)). PageRank adds a perturbation to the Invariant method in order to deal with irreducible problems.

As mentioned in the introduction, some scoring functions seem more appropriate for some scoring problems than for others, depending on their respective interpretations. For example, we will argue that the Invariant scoring function is appealing when the scoring problem is interpreted as a web page ranking problem, while the fair-bets scoring function makes more sense when one deals with different kinds of tournaments.

Consider the Invariant scoring function. If the scoring problem represents a web page ranking problem, then the Invariant score v_i of web page i is the weighted sum of the reference shares that web page i gets from each of the web pages, the weights being the scores of the respective web pages.

Consider now the fair-bets scoring function. When the scoring problem represents a tournament, player i's score is interpreted as the value of a win over i. The fair-bets scores are chosen so that for each player, the total value of its wins equals the value of its losses.

Note that the Invariant scores are chosen so that the score of web page i can be written as a function of the citations it gets from the other web pages, and not of the references it provides. This seems reasonable in the context of web page ranking problems, because making references is "free": it does not reduce the number of citations one gets. On the other hand, the fair-bets scores are chosen so as to balance the values of the wins and

[5] That every irreducible stochastic matrix have a unique stationary distribution is a standard result in finite Markov Chain theory. See Kemeny and Snell (1976)

the losses of each player. In a tournament, winning and loosing are mutually exclusive outcomes. If you don't win, you lose.

Example 1 Consider the scoring problem $S = \langle N, A \rangle$ where $N = \{a, b, c, d\}$ and

$$A = \begin{array}{c} \\ a \\ b \\ c \\ d \end{array} \begin{array}{c} \begin{matrix} a & b & c & d \end{matrix} \\ \begin{pmatrix} 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \end{array}.$$

The corresponding graph, ignoring the numerical labels, is shown in Figure 1.

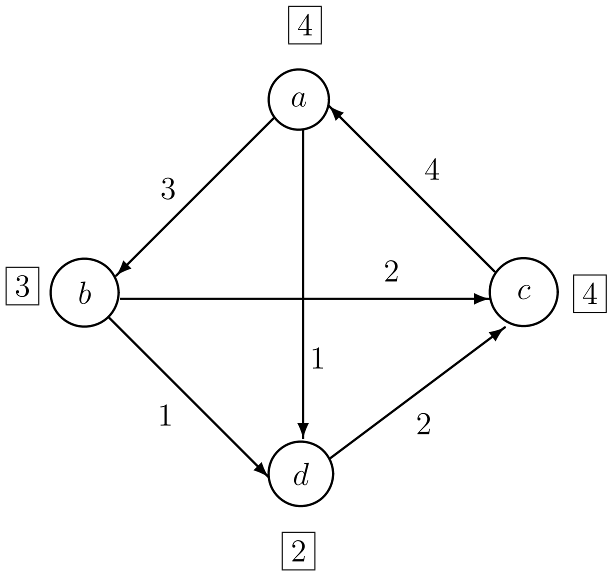

Figure 1: A scoring problem.

The numbers labeling the edges pointing to node i, represent i's fair-bets score. For instance, $F(c) = 2$ and $F(d) = 1$. Note that for all i, the sum of the numbers labeling the incoming edges equals the sum of the numbers labeling the outgoing edges. These common values appear framed next to the nodes and they represent each node's Invariant score. For instance, $I(c) = 4$ and $I(d) = 2$.

3.2 Intuitive interpretation of the two methods

In this subsection we provide two interpretations of the Invariant and fair-bets methods. The first interpretation is based on simple Markov processes, whereas the other is based on graph-theoretic concepts.

In the context of web browsing, consider a "random surfer" who keeps clicking on the pages' hyperlinks uniformly at random. Then, the long run probability that this random surfer is at any particular page is precisely this page's invariant score. Equivalently, a page's invariant score is the proportion of time a random surfer visits that page (see Brin and Page (1998)).

For the fair-bets scores, the interpretation is based on the following "ping-pong" protocol (see Laslier (1997)). A pair of players play against each other and the odds are given by the ratio a_{ij}/a_{ji}. The player who loses cedes his position at the table, and the winner, say player i, plays the next game with an opponent that is selected randomly according to the proportions of games played against i. That is, opponent j is selected with probability $\frac{a_{ij}+a_{ji}}{a_{i*}+a_{*i}}$. The process continues ad infinitum. This defines a Markov process whose stochastic matrix is $(A+\mathrm{diag}(a_{i*}))(\mathrm{diag}(a_{i*}+a_{*i}))^{-1}$, where the ij-entry is the probability that the next game is played by i, given that player j is currently playing. The stationary distribution of this Markov process gives the proportion of time that each player remains at the table. Alternatively, it is the vector of long run probabilities that each player is the winner. This long run probability has two components. One is related to the chances of winning once the opponent is selected, and the other component is related to the chances of being selected, which is itself related to the number of games that the player plays against other players as given by the matrix A. Note that not all players play the same number of games according to the matrix A: player i plays $a_{i*} + a_{*i}$ games. Therefore, if we are interested only in the winning probabilities, we should divide the stationary probabilities of the above Markov chain by the corresponding number of games. It turns out that the resulting vector is proportional to the fair-bets scores. Thus, the fair bets score of player i is his stationary probability of being a winner, divided by the number of games he plays, properly normalized so that the sum of the scores is one. To see this, let v be the fair-bets scores of $\langle N, A\rangle$ and consider the matrix $A + \text{diag}(a_{i*})$. Then,

$$\begin{aligned}(A + \text{diag}(a_{i*}))v &= Av + \text{diag}(a_{i*})v \\ &= C_A v + \text{diag}(a_{i*})v \\ &= \text{diag}(a_{*i} + a_{i*})v.\end{aligned}$$

But since $(A+\text{diag}(a_{i*}))v = (A+\text{diag}(a_{i*}))\text{diag}^{-1}(a_{*i}+a_{i*})\text{diag}(a_{*i}+a_{i*})v$, we obtain that $\text{diag}(a_{*i}+a_{i*})v$ is proportional to the stationary distribution of $(A+\text{diag}(a_{i*}))\text{diag}^{-1}(a_{*i}+a_{i*})$. Thus, v can be viewed as the "per game" stationary distribution.

We now turn to the graph-theoretic interpretation.[6] Let $S = \langle N, A\rangle$ be a scoring problem. A directed graph over vertex set N is a set of directed edges $E \subseteq N \times N$. An *i-tree* is a directed graph over N, such that for every vertex $j \in N \setminus \{i\}$, there is a unique path from i to j. For each $i \in N$, let $\mathcal{T}_i$ denote the set of all i-trees. An *i-cycle* is a directed graph C that consists of an i-tree with an additional edge (j, i) for some $j \in N \setminus \{i\}$. In other words, an i-cycle is a set of edges $C = T \cup \{(j, i)\}$, where T is an i-tree and $j \in N \setminus \{i\}$. Denoting by $\mathcal{C}_i$ the set of all i-cycles, we have

$$\mathcal{C}_i = \{T_i \cup \{(j, i)\} : T_i \in \mathcal{T}_i,\ j \neq i\}. \tag{5}$$

Note that if C is an i-cycle, then there is $j \neq i$ such that $C \setminus \{(i, j)\}$ is a j-tree. Consequently, C can be written as $C = T_j \cup \{(i, j)\}$ for some $j \neq i$ and hence it is a j-cycle. Conversely, if $C = T_j \cup \{(i, j)\}$ for some $j \neq i$, then C is also an i-cycle. As a result we can write

$$\mathcal{C}_i = \{T_j \cup \{(i, j)\} : T_j \in \mathcal{T}_j,\ j \neq i\}. \tag{6}$$

For any set of directed edges E define

$$P(E) = \prod_{(i,j)\in E} a_{ij}$$

[6] This interpretation is based on Lemma 3.1 in Freidlin and Wentzell (1998).

and for each $i \in N$, define the following two values:

$$\begin{aligned}\nu(i) &= \sum_{T\in\mathcal{T}_i} P(T)\\ \mu(i) &= \sum_{C\in\mathcal{C}_i} P(C).\end{aligned}$$

In the case that S is a tournament, $\nu(i)$ and $\mu(i)$ are the number of i-trees and i-cycles in the corresponding tournament graph.

It turns out that $(\nu(i))_{i\in N}$ solves (3) and $(\mu(i))_{i\in N}$ solves (1). That is, they are proportional to the fair-bets scores and to the Invariant scores, respectively. To see this, note that, using the equality (5),

$$\begin{aligned}\sum_{C\in\mathcal{C}_i} P(C) &= \sum_{T\in\mathcal{T}_i}\sum_{j\neq i} P(T)a_{ji}\\ &= \sum_{j\neq i} a_{ji}\sum_{T\in\mathcal{T}_i} P(T)\\ &= \sum_{j\neq i} a_{ji}\nu(i).\end{aligned}$$

Similarly, using (6)

$$\begin{aligned}\sum_{C\in\mathcal{C}_i} P(C) &= \sum_{T\in\mathcal{T}_j}\sum_{j\neq i} P(T)a_{ij}\\ &= \sum_{j\neq i} a_{ij}\sum_{T\in\mathcal{T}_j} P(T)\\ &= \sum_{j\neq i} a_{ij}\nu(j).\end{aligned}$$

Consequently,

$$\sum_{j\neq i}\nu(i)a_{ji} = \sum_{j\neq i}\nu(j)a_{ij}.$$

The last equality shows that $(\nu(i))_{i\in N}$ solves (3). But then, it follows that $(\sum_{j\neq i}\nu(i)a_{ji})_{i\in N}$ solves (1). And since $\mu(i) = \sum_{C\in\mathcal{C}_i} P(T) = \sum_{j\neq i}\nu(i)a_{ji}$ we have that $(\mu(i))_{i\in N}$ solves (1). It follows that, in particular, when $S = \langle N, A\rangle$ is a tournament, the fair-bets scoring function ranks the nodes of S by the number of spanning trees rooted at them, whereas the Invariant scoring function ranks the nodes by the number of r-cycles rooted at nodes $r \in N$.

3.3 Properties of scoring functions

In this section we state several simple properties of scoring functions that one might reasonably require. We may find some properties to be more appealing than others depending on the application at hand. These properties will be used in our axiomatizations in the remainder of the paper.

A scoring problem $S = \langle N, A \rangle$ is *balanced* if for all $i \in N, a_{i*} = a_{*i}$. It is *regular* if in addition, for all $i, j \in N, a_{i*} = a_{j*}$.

In the context of web page ranking problems, a balanced problem is one where each web page's number of citations equals its number of references. If the scoring problem represents a tournament, balancedness of a problem means that for each player, the number of victories equals the number of losses. If, in addition, for all players the number of victories (and losses) are the same, then the problem is regular.

Definition 1 A scoring function, f, is *uniform* if for all regular problems $S = \langle N, A \rangle$, we have $f(S) = (1/|N|)_{i \in N}$.

Uniformity of a scoring function requires that if all players (web pages) have the same number of wins (citations) and the same number of losses (references), then they should be all equally ranked. For the record, we state without proof the following lemma:

Lemma 1 Both the fair-bets and the Invariant scoring functions are uniform.

One can strengthen uniformity in the following way:

Definition 2 A scoring function, f, is *strongly uniform* if for all balanced problems $S = \langle N, A \rangle$, we have $f(S) = (1/|N|)_{i \in N}$.

The strong uniformity of a scoring function is more appealing when the scoring problem represents a tournament. It says that if each player's wins/losses ratio is 1, then they should all be ranked equally.

Lemma 2 The fair-bets scoring function is strongly uniform.

Proof : Let $S = \langle N, A \rangle$ be a balanced scoring problem. Its fair-bets scores $(v_i)_{i \in N}$ are the only ones that satisfy

$$\sum_j a_{ji} v_i = \sum_j a_{ij} v_j \qquad \text{for all } i \in N.$$

But since S is balanced, we have $\sum_j a_{ji} = \sum_j a_{ij}$ for all $i \in N$, and hence the vector $v = (1/|N|)_{i \in N}$ solves the above equation. □

Definition 3 A scoring function, f, is *neutral* if for all regular scoring problems $\langle N, A \rangle$ and for all symmetric matrices B whose main diagonal entries are all 0, and for which $\langle N, A+B \rangle$ is also a scoring problem, if $f(\langle N, A \rangle) = (1/|N|)_{i \in N}$, then $f(\langle N, A+B \rangle) = (1/|N|)_{i \in N}$.

Neutrality is an appealing property in the context of generalized tournaments. It says that if the players in a regular scoring problem are ranked equally, then if we add some more games, and each player wins half of the additional games he plays and losses the other half, then the players should remain equally ranked.

The next lemma shows the relationship between neutrality, uniformity and strong uniformity.

Lemma 3 The scoring function f is uniform and neutral if and only if f is strongly uniform.

Proof : It is clear that strong uniformity implies uniformity. Also, since by adding a symmetric matrix to the matrix of a regular problem one gets a balanced problem, strong uniformity implies neutrality. To show the other direction, let f be a uniform and neutral scoring function and let $S = \langle N, A \rangle$ be a balanced scoring problem with $|N| > 2$ (if $|N| \le 2$, then S is regular and there is nothing to prove). Let $b = \max\{a_{ij} : i, j \in N\}$ and define the non-negative and symmetric matrix $B = (b_{ij})$ as follows:

$$b_{ij} = \begin{cases} 0 & \text{if } i = j \\ b - \dfrac{a_{ij} + a_{ji}}{2} & \text{otherwise.} \end{cases}$$

The scoring problem $S' = \langle N, A + B \rangle$ is regular. Indeed, for any $i \in N$

$$\begin{aligned} \sum_{j \ne i} (a_{ji} + b_{ji}) = \sum_{j \ne i} (a_{ij} + b_{ij}) &= \sum_{j \ne i} (a_{ij} + b - \frac{a_{ij} + a_{ji}}{2}) \\ &= \sum_{j \ne i} (b + \frac{a_{ij} - a_{ji}}{2}) \\ &= (n-1)b. \end{aligned}$$

By uniformity of f we have $f(\langle N, A + B \rangle) = (1/|N|)_{i \in N}$. But since $A = (A + B) + (-B)$, by neutrality of f, $f(\langle N, A \rangle) = (1/|N|)_{i \in N}$. □

As a result we have the following:

Corollary 1 The fair-bets scoring function is neutral.

Proof : An immediate consequence of Lemma 2 and Lemma 3. □

The following property is appealing in the context of web page ranking problems. It says that if web pages are ranked in proportion to their number of citations in a regular problem, they should still be ranked according to their number of citations after some pairs of web pages exchange equal number of mutual extra links. Formally,

Definition 4 A scoring function, f, is *weakly additive* if for all regular scoring problems $\langle N, A \rangle$ and for all symmetric matrices B whose main diagonal entries are all 0, and for which $\langle N, A + B \rangle$ is also a scoring problem, if $f(\langle N, A \rangle) \propto (a_{i*})_{i \in N}$, then $f(\langle N, A + B \rangle) \propto (a_{i*} + b_{i*})_{i \in N}$.

Lemma 4 The Invariant scoring function is weakly additive.

Proof : Let $S = \langle N, A \rangle$ be a regular scoring problem such that $I(S) \propto (a_{i*})_{i \in N}$, and let B be an $|N| \times |N|$ matrix, with 0 as its main diagonal entries, such that $S' = \langle N, A + B \rangle$ is a scoring problem. Since S is regular, by Lemma 1, $F(S) = (1/|N|)_{i \in N}$, and since F satisfies neutrality, $F(S') = (1/|N|)_{i \in N}$ and as a result $I(S') \propto C_{A+B}F(S') = (C_A + C_B)(1/|N|)_{i \in N} \propto (a_{*j} + b_{*j})_{j \in N}$. □

Definition 5 A scoring function, f, is *invariant to reference intensity* if for all scoring problems $S = \langle N, A \rangle$ and for all $|N| \times |N|$ diagonal matrices Λ with positive diagonal entries, we have

$$f(N, A) = f(N, A\Lambda).$$

This property makes sense if the scoring problem represents a web page ranking problem. It says that if the references of a web page are multiplied by a positive number, the scores and resulting ranking should not change. In other words, a web page should not be able to affect the ranking simply by multiplying its references by a constant.

An analogous property that is more relevant in the context of tournaments, is the following.

Definition 6 A scoring function, f, is *inversely proportional to losses* if for all balanced scoring problems $\langle N, A \rangle$ such that $f(N, A) = (1/|N|)_{i \in N}$, and for all $(\lambda_i)_{i \in N} \gg 0$ we have

$$\frac{f_i(\langle N, A\,\mathrm{diag}(\lambda_i)_{i \in N} \rangle)}{f_j(\langle N, A\,\mathrm{diag}(\lambda_i)_{i \in N} \rangle)} = \frac{\lambda_j}{\lambda_i} \qquad \text{for all } i, j \in N,$$

or in matrix notation,

$$f(\langle N, A\,\mathrm{diag}(\lambda_i)_{i \in N} \rangle) \propto (\frac{1}{\lambda_i})_{i \in N}.$$

This property requires that if in a given balanced problem all players are equally ranked, and if one player's losses are multiplied by a constant, its relative score should be divided by that constant.

Lemma 5 The Invariant scoring function is invariant to reference intensity.

Proof : Let $S = \langle N, A \rangle$ and $S' = \langle N, A\Lambda \rangle$ be two scoring problems where $\Lambda = \mathrm{diag}(\lambda_i)_{i \in N}$ for some positive vector $(\lambda_i)_{i \in N}$. Let $(v'_i)_{i \in N} = I(S')$. Then, by definition

$$v'_i = \sum_j \frac{\lambda_j a_{ij}}{\lambda_j a_{*j}} v'_j \qquad \text{for all } i \in N,$$

or equivalently,

$$v'_i = \sum_j \frac{a_{ij}}{a_{*j}} v'_j \qquad \text{for all } i \in N.$$

But then $(v'_i)_{i \in N} = I(S)$, which proves that I is invariant to reference intensity. □

Lemma 6 The fair-bets scoring function is inversely proportional to losses.

Proof : We will prove a stronger claim. Namely, that the player's relative scores awarded by the fair-bets scoring function are homogeneous of degree -1 in the player's column. Let $S = \langle N, A \rangle$ be a scoring problem and let $\Lambda = \mathrm{diag}(\lambda_i)_{i \in N}$ be a diagonal matrix with positive main diagonal. Since I is invariant to reference intensity, we have that $I(\langle N, A\Lambda \rangle) = I(\langle N, A \rangle)$ which holds if and only if $C_{A\Lambda} F(\langle N, A\Lambda \rangle) \propto C_A F(\langle N, A \rangle)$, which holds if and only if $F(\langle N, A\Lambda \rangle) \propto \Lambda^{-1} F(\langle N, A \rangle)$, where we use the fact that $C_{A\Lambda} = C_A \Lambda$. Therefore, F is inversely proportional to losses. □

4 The main results

4.1 Two dual axiomatizations

We can now state our first result.

Theorem 1 Let $N \subseteq \mathbb{N}$ be a finite set of at least two players. The fair-bets is the only scoring function defined on $\mathcal{S}_0(N)$ that satisfies uniformity, neutrality and inverse proportionality to losses.

Proof : It was already shown that the fair-bets scoring function satisfies the three properties. In order to show that it is the only scoring function that does so let f be a scoring function that satisfies the three properties and let $S = \langle N, A \rangle$ be a scoring problem. Let $v = (v_i)_{i \in N} = F(S)$. We need to show that $f(S) = v$. Consider the auxiliary scoring problem $S' = \langle N, A\text{diag}(v) \rangle$. Since $(v_i)_{i \in N}$ are the fair-bets scores of S, we have $\sum_j a_{ji} v_i = \sum_j a_{ij} v_j$ for all $i \in N$, which means that S' is a balanced scoring problem. Consequently, since by Lemma 3 f is strongly uniform, $f(S') = (1/|N|, \ldots, 1/|N|)$. Since f is inversely proportional to losses, we have that

$$f(S) = f(N, A\,\text{diag}(v)(\text{diag}(v))^{-1}) \propto (v_i)_{i \in N} = F(S). \tag{7}$$

Since $f(S)$ and $F(S)$ are in Δ_N, $f(S) = F(S)$. □

Corollary 2 Let $N \subseteq I\!N$ be a finite set of at least two players. The Invariant function is the only scoring function defined on $\mathcal{S}_0(N)$ that is uniform, weakly additive and invariant to reference intensity.

Proof : We have already proved that the Invariant scoring function satisfies the three properties. Therefore, we need to show that it is the only one that does so. If two scoring functions f_1 and f_2 are uniform, weakly additive and invariant to reference intensity, then the scoring functions g_1 and g_2 defined by

$$g_i(\langle N, A \rangle) = \frac{C_A^{-1} f_i(\langle N, A \rangle)}{\|C_A^{-1} f_i(\langle N, A \rangle)\|} \qquad i = 1, 2$$

both satisfy uniformity, neutrality and inverse proportionality to losses. By Theorem 1, g_1 and g_2 are the same function, which implies that f_1 and f_2 are the same. □

4.2 Independence of the axioms

Next we show that the properties that characterize the Invariant scoring function in Corollary 2 are logically independent, namely no two properties imply the third one.

1. Consider the scoring function f that assigns to each scoring problem $S = \langle N, A\rangle$ the scores $f(S) \in \Delta_N$ such that $f(S) \propto (a_{i*})_{i\in N}$. It is easily seen that this function satisfies uniformity and weak additivity. It does not satisfy invariance to reference intensity.

2. Consider the function g that assigns to each scoring problem $S = \langle N, A\rangle$ the scores $g(S) \in \Delta_N$ such that $g(S) \propto (\sum_{j\in N} \frac{a_{ij}}{a_{*j}})_{i\in N}$. It is easily checked that g satisfies uniformity and invariance to reference intensity. Since $g \neq I$, it follows from Corollary 2 that g does not satisfy weak additivity.

3. Consider the function h that assigns to each scoring problem $S = \langle N, A\rangle$ the scores $h(S) \in \Delta_N$ such that $h(S) \propto (i \sum_{j\in N} \frac{a_{ij}}{a_{*j}})_{i\in N}$. It can be seen that h satisfies weak additivity (trivially) and invariance to reference intensity, but it does not satisfy uniformity.

The fact that the three properties that we used in Theorem 1 to characterize the fair-bets scoring function are independent, can be easily checked using the functions φ, ϕ and ψ defined as follows: $\varphi(S) = \frac{f(S)C_A^{-1}}{\|f(S)C_A^{-1}\|}$, $\phi(S) = \frac{g(S)C_A^{-1}}{\|g(S)C_A^{-1}\|}$, and $\psi(S) = \frac{h(S)C_A^{-1}}{\|h(S)C_A^{-1}\|}$.

4.3 An alternative axiomatization

Note that in the characterizations presented in Section 4.1, the class of problems can be taken as the set of all scoring problems in $\mathcal{S}_0$ with a fixed set of agents. In the following characterization, a variable number of agents is needed. More precisely, the class of problems should be such that whenever it contains all the scoring problems in $\mathcal{S}_0$ with a fixed set of agents, N, it also contains all the problems with any nonempty subset of N.

In the context of journal citations, Palacios-Huerta and Volij (2002) characterize the Invariant function using the axioms of weak homogeneity and weak consistency, along with invariance to reference intensity. By strengthening the weak homogeneity and consistency axioms, we can present an alternative characterization of the fair-bets scoring function.

Definition 7 A scoring function, f, satisfies *reciprocity*, if for all two-player scoring problems $S = \langle \{i, j\}, A\rangle$, $f_i(S)/f_j(S) = a_{ij}/a_{ji}$.

Reciprocity requires that in two-player problems, the relative scores of the players be their relative number of victories. [7] It is a very weak requirement.

In order to define consistency, we need the following definition. Let $S = \langle N, A\rangle \in \mathcal{S}_0$ be a scoring problem with at least three players and let k be one of the players. *The reduced scoring problem with respect to* k is given by $S^k = \langle N \setminus \{k\}, A^k\rangle$ where $A^k = (a^k_{ij})$ is defined by

$$a^k_{ij} = \begin{cases} 0 & \text{if } i = j \\ a_{ij} + \dfrac{a_{ik}a_{kj}}{\sum_{t\in N} a_{tk}} & \text{otherwise.} \end{cases}$$

The reduced problem with respect to k consists of all the players in the original scoring problem, except for k, and a modified $(|N| - 1) \times (|N| - 1)$ matrix. It summarizes all the inter-relations between the players of the original problem except for player k, including their indirect relationships through player k. The idea is to describe these relationships using a matrix that, though smaller than the one in the original scoring problem, still contains all the relevant information. In order to account for the influence of k, we add the term $\frac{a_{ik}a_{kj}}{\sum_{t\in N} a_{tk}}$ to each of the entries a_{ij}, $i, j \neq k$, $i \neq j$, of the original matrix.

In the context of generalized tournaments the entry a_{ij} represents the number of times that j was defeated by i. If we delete k's row, and in particular if we delete the entry a_{kj}, we omit the information about the "indirect victories" of each of the players over j via k. To compensate for this omission, we take the entry a_{kj} and distribute it among the players in proportion to their respective number of victories over k. In particular, if player i recorded no victories over player k, he receives no share of this distribution.

The reduced problem with respect to player k intends to reflect the relationships of the players in $N \setminus \{k\}$ in the original scoring problem. If it does so, it would be natural to require from a scoring function to assign the players the same relative scores in both problems. This is precisely the consistency requirement that we define next.

Definition 8 The scoring function f satisfies *consistency* if for all scoring problems $S = \langle N, A\rangle$ with at least three players and for all $k \in N$,

$$\frac{f_i(S^k)}{f_j(S^k)} = \frac{f_i(S)}{f_j(S)} \qquad \text{for all } i, j \neq k.$$

[7] Note that since A is irreducible, this ratio is well-defined.

The consistency principle has been extensively applied in the axiomatic literature. See Thomson (2000) for a comprehensive survey on consistency.

Theorem 2 The fair-bets scoring function is the unique scoring function on $\mathcal{S}_0$ that satisfies reciprocity and consistency.

Proof : We show first, by induction, that there cannot be two functions that satisfy the two properties. Suppose that both f and g satisfy the two properties and coincide for all n-player scoring problems. Since f and g satisfy reciprocity, this is true for $n = 2$. Let $S = \langle N, A\rangle$ be an $(n+1)$-player problem. Let $k \in N$ and let $i, j \in N \setminus \{k\}$. By consistency of f and g, we have

$$\frac{f_i(S)}{f_j(S)} = \frac{f_i(S^k)}{f_j(S^k)} = \frac{g_i(S^k)}{g_j(S^k)} = \frac{g_i(S)}{g_j(S)}.$$

Since i and j were chosen arbitrarily, this implies that $f(S) = g(S)$.

It remains to show that the fair-bets scoring function satisfies both axioms. That it satisfies reciprocity is easy to check and is left to the reader. To see that it satisfies consistency, let $S = \langle N, A\rangle$ be a scoring problem with at least three players and let $(v_i)_{i\in N}$ be the corresponding fair-bets scores. That is,

$$\sum_{j\in N} a_{ji}v_i = \sum_{j\in N} a_{ij}v_j \qquad \text{for all } i \in N. \tag{8}$$

Let $k \in N$ and consider the reduced scoring problem S^k. We need to show that $F(S^k) \propto (v_i)_{i\in N\setminus\{k\}}$. It is enough to show that

$$\sum_{j\in N\setminus\{k\}} a^k_{ji}v_i = \sum_{j\in N\setminus\{k\}} a^k_{ij}v_j \qquad \text{for all } i \in N \setminus \{k\}.$$

First note that

$$\begin{aligned}\sum_{j\in N\setminus\{k\}} a^k_{ji}v_i &= \sum_{j\in N\setminus\{i,k\}} (a_{ji} + \frac{a_{jk}a_{ki}}{\sum_{t\in N} a_{tk}})v_i \\ &= \sum_{j\in N\setminus\{k\}} (a_{ji} + \frac{a_{jk}a_{ki}}{\sum_{t\in N} a_{tk}})v_i - \frac{a_{ik}a_{ki}}{\sum_{t\in N} a_{tk}}v_i\end{aligned}$$

$$
\begin{aligned}
&= \sum_{j\in N\setminus\{k\}} a_{ji}v_i + a_{ki}\sum_{j\in N}\frac{a_{jk}}{\sum_{t\in N} a_{tk}}v_i - \frac{a_{ik}a_{ki}}{\sum_{t\in N} a_{tk}}v_i \\
&= \sum_{j\in N} a_{ji}v_i - \frac{a_{ik}a_{ki}}{\sum_{t\in N} a_{tk}}v_i.
\end{aligned}
$$

Next,

$$
\begin{aligned}
\sum_{j\in N\setminus\{k\}} a^k_{ij}v_j &= \sum_{j\in N\setminus\{i,k\}} (a_{ij} + \frac{a_{ik}a_{kj}}{\sum_{t\in N} a_{tk}})v_j \\
&= \sum_{j\in N\setminus\{k\}} (a_{ij} + \frac{a_{ik}a_{kj}}{\sum_{t\in N} a_{tk}})v_j - \frac{a_{ik}a_{ki}}{\sum_{t\in N} a_{tk}}v_i \\
&= \sum_{j\in N\setminus\{k\}} a_{ij}v_j + a_{ik}\frac{\sum_{j\in N} a_{kj}v_j}{\sum_{t\in N} a_{tk}} - \frac{a_{ik}a_{ki}}{\sum_{t\in N} a_{tk}}v_i \\
&= \sum_{j\in N\setminus\{k\}} a_{ij}v_j + a_{ik}\frac{\sum_{j\in N} a_{jk}v_k}{\sum_{t\in N} a_{tk}} - \frac{a_{ik}a_{ki}}{\sum_{t\in N} a_{tk}}v_i \\
&= \sum_{j\in N} a_{ij}v_j - \frac{a_{ik}a_{ki}}{\sum_{t\in N} a_{tk}}v_i.
\end{aligned}
$$

The previous to last equality follows from the fact that, by equation (8), $\sum_{j\in N\setminus\{k\}} a_{kj}v_j = \sum_{j\in N\setminus\{k\}} a_{jk}v_k$. Using equation (8) again, we get the desired result. □

Remarks:

1. The class of problems can be restricted to $\cup_{N'\subseteq N}\mathcal{S}_0(N')$ for some $N' \in \mathcal{N}$.

2. The axioms used in the above characterization are logically independent: the Invariant scoring function satisfies reciprocity but does not satisfy consistency; and the scoring function that assigns to every scoring problem $\langle N, A\rangle$ the vector $(1/|N|)_{i\in N}$ satisfies consistency but does not satisfy reciprocity.

References

Arrow, K. J. (1963). *Social Choice and Individual Values* (2nd. ed.). New Haven: Yale University Press.

Brin, S. and L. Page (1998). The anatomy of large-scale hypertextual web search engine. *Computer Networks 30(1-7)*, 107–117.

Chakrabarti, S., B. Dom, D. Gibson, J. Kleinberg, R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins (1999). Hypersearching the web. *Scientific American* (June issue).

Conner, G. R. and C. P. Grant (2000). An extension of Zermelo's model for ranking by paired comparisons. *European Journal of Applied Mathematics 11*, 225–247.

Daniels, H. E. (1969). Round-robin tournament scores. *Biometrika 56*, 295–299.

David, H. A. (1987). Ranking from unbalanced paired-comparison data. *Biometrika 74*, 432–436.

David, H. A. (1988). *The Method of Paired Comparisons* (2 ed.). London: Charles Griffin and Company.

Elo, A. E. (1978). *The Rating of Chess Players, Past and Present.* Arco Publishing Inc.

Freidlin, M. I. and A. D. Wentzell (1998). *Random Perturbations of Dynamical Systems* (2nd ed.). Springer.

Henriet, D. (1985). The Copeland choice function: an axiomatic characterization. *Social Choice and Welfare 2*, 49–63.

Herings, J.-J., G. van der Laan, and D. Talman (2001, October). Measuring the power of nodes in digraphs. Discussion paper.

Kano, M. and A. Sakamoto (1985). Ranking the vertices of a paired comparison digraph. *SIAM Journal on Algebraic and Discrete Methods 6*, 79–92.

Keener, J. P. (1993). The Perron-Frobenius theorem and the ranking of footbal teams. *SIAM Review 35*, 80–93.

Kemeny, J. G. and J. L. Snell (1976). *Finite Markov Chains.* Berlin: Springer-Verlag.

Kendall, M. G. (1955). Further contributions to the theory of paired comparisons. *Biometrics 11*, 43–62.

Kleinberg, J. (1999). Authoritative sources in a hyperlinked environment. *Journal of the ACM 46*, 604–632.

Laslier, J. (1997). *Tournament Solutions and Majority Voting*. Berlin: Springer.

Levchenkov, V. S. (1995). Self-consistent rule for group choice. I: Axiomatic approach. Discussion Paper 95–3, Conservatoire National des Arts et Métiers.

Merlin, V. and D. G. Saari (1996). Copeland method I: Dictionaries and relationships. *Economic Theory 8*, 51–76.

Moon, J. W. (1968). *Topics on Tournaments*. New York: Holt, Rinehart and Winston.

Moon, J. W. and N. J. Pullman (1970). On generalized tournament matrices. *SIAM Review 12*, 384–399.

Moulin, H. (1988). *Axioms of Cooperative Decision Making*. Cambridge, U. K.: Cambridge University Press.

Page, L., S. Brin, R. Motwani, and T. Winograd (1999). The PageRank citation ranking: Bringing order to the web. Technical report, Stanford University.

Palacios-Huerta, I. and O. Volij (2002). The measurement of intellectual influence. mimeo.

Pinski, G. and F. Narin (1976). Citation influence for journal aggregates of scientific publications: Theory, with applications to the literature of physics. *12*, 297–312.

Rubinstein, A. (1980). Ranking the participants in a tournament. *SIAM Journal of Applied Mathematics 38*, 108–111.

Saari, D. G. (2001). *Decisions and Elections: Explaining the Unexpected*. Cambridge University Press.

Saaty, T. L. (1980). *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. McGraw Hill.

Sen, A. (1970). *Individual Choice and Social Welfare*. Holden Day.

Thomson, W. (2000). Consistent allocation rules. Mimeo, University of Rochester.

van den Brink and R. Gilles (2000). Measuring domination in directed networks. *Social Networks 22*, 141–157.

Wei, T. (1952). The algebraic foundation of ranking theory. Ph. D. Thesis, Cambridge University.

Zermelo, E. (1929). Die Berechnung der Turnier-Ergebnisse als ein Maximalproblem der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift 29*, 436–460.