

On the Axiomatic Foundations of Ranking Systems

Alon Altman and Moshe Tennenholtz
Faculty of Industrial Engineering and Management
Technion – Israel Institute of Technology
Haifa 32000
Israel

Abstract

Reasoning about agent preferences on a set of alternatives, and the aggregation of such preferences into some social ranking is a fundamental issue in reasoning about multi-agent systems. When the set of agents and the set of alternatives coincide, we get the ranking systems setting. A famous type of ranking systems are page ranking systems in the context of search engines. In this paper we present an extensive axiomatic study of ranking systems. In particular, we consider two fundamental axioms: Transitivity, and Ranked Independence of Irrelevant Alternatives. Surprisingly, we find that there is no general social ranking rule that satisfies both requirements. Furthermore, we show that our impossibility result holds under various restrictions on the class of ranking problems considered. Each of these axioms can be individually satisfied. Moreover, we show a complete axiomatization of approval voting using one of these axioms.

1 Introduction

The ranking of agents based on other agents' input is fundamental to multi-agent systems (see e.g. [Resnick *et al.*, 2000]). Moreover, it has become a central ingredient of a variety of Internet sites, where perhaps the most famous examples are Google's PageRank algorithm[Page *et al.*, 1998] and eBay's reputation system[Resnick and Zeckhauser, 2001].

This basic problem introduces a new social choice model. In the classical theory of social choice, as manifested by Arrow[1963], a set of agents/voters is called to rank a set of alternatives. Given the agents' input, i.e. the agents' individual rankings, a social ranking of the alternatives is generated. The theory studies desired properties of the aggregation of agents' rankings into a social ranking. In particular, Arrow's celebrated impossibility theorem[Arrow, 1963] shows that there is no aggregation rule that satisfies some minimal requirements, while by relaxing any of these requirements appropriate social aggregation rules can be defined. The novel feature of the ranking systems setting is that the set of agents and the set of alternatives **coincide**. Therefore, in such setting one may need to consider the transitive effects of voting. For example, if agent a reports on the importance of (i.e. votes

for) agent b then this may influence the credibility of a report by b on the importance of agent c ; these indirect effects should be considered when we wish to aggregate the information provided by the agents into a social ranking.

Notice that a natural interpretation/application of this setting is the ranking of Internet pages. In this case, the set of agents represents the set of Internet pages, and the links from a page p to a set of pages Q can be viewed as a two-level ranking where agents in Q are preferred by agent(page) p to the agents(pages) which are not in Q . The problem of finding an appropriate social ranking in this case is in fact the problem of (global) page ranking. Particular approaches for obtaining a useful page ranking have been implemented by search engines such as Google[Page *et al.*, 1998].

The theory of social choice consists of two complementary axiomatic perspectives:

- The descriptive perspective: given a particular rule r for the aggregation of individual rankings into a social ranking, find a set of axioms that are sound and complete for r . That is, find a set of requirements that r satisfies; moreover, every social aggregation rule that satisfies these requirements should coincide with r . A result showing such an axiomatization is termed a *representation theorem* and it captures the exact essence of (and assumptions behind) the use of the particular rule.
- The normative perspective: devise a set of requirements that a social aggregation rule should satisfy, and try to find whether there is a social aggregation rule that satisfies these requirements.

Many efforts have been invested in the descriptive approach in the framework of the classical theory of social choice. In that setting, representation theorems have been presented to major voting rules such as the majority rule[May, 1952]. Recently, we have successfully applied the descriptive perspective in the context of ranking systems by providing a representation theorem[Altman and Tennenholtz, 2005] for the well-known PageRank algorithm [Page *et al.*, 1998], which is the basis of Google's search technology.

An excellent example for the normative perspective is Arrow's impossibility theorem mentioned above. In [Tennenholtz, 2004], we presented some preliminary results for ranking systems where the set of voters and the set of alternatives coincide. However, the axioms presented in that work consist

of several very strong requirements which naturally lead to an impossibility result.

In this paper we provide an extensive study of ranking systems. We introduce two fundamental axioms. One of these axioms captures the transitive effects of voting in ranking systems, and the other adapts Arrow’s well-known independence of irrelevant alternatives (IIA) axiom to the context of ranking systems. Surprisingly, we find that no general ranking system can simultaneously satisfy these two axioms! We further show that our impossibility result holds under various restrictions on the class of ranking problems considered. On the other hand, we show that each of these axioms can be individually satisfied. Moreover, we use our IIA axiom to present a positive result in the form of a representation theorem for the well-known approval voting ranking system, which ranks the agents based on the number of votes received. This axiomatization shows that when ignoring transitive effects, there is only one ranking system that satisfies our IIA axiom.

This paper is structured as follows: Section 2 formally defines our setting and the notion of ranking systems. Sections 3 and 4 introduce our axioms of Transitivity and Ranked Independence of Irrelevant Alternatives respectively. Our main impossibility result is presented in Section 5, and further strengthened in Section 6. Our positive result, in the form of an axiomatization for the Approval Voting ranking system is presented in Section 7. Finally, some concluding remarks are given in Section 8.

2 Ranking Systems

Before describing our results regarding ranking systems, we must first formally define what we mean by the words “ranking system” in terms of graphs and linear orderings:

Definition 2.1. Let A be some set. A relation $R \subseteq A \times A$ is called an *ordering* on A if it is reflexive, transitive, and complete. Let $L(A)$ denote the set of orderings on A .

Notation 2.2. Let \preceq be an ordering, then \simeq is the equality predicate of \preceq , and \prec is the strict order induced by \preceq . Formally, $a \simeq b$ if and only if $a \preceq b$ and $b \preceq a$; and $a \prec b$ if and only if $a \preceq b$ but not $b \preceq a$.

Given the above we can define what a ranking system is:

Definition 2.3. Let \mathbb{G}_V be the set of all graphs with vertex set V . A *ranking system* F is a functional that for every finite vertex set V maps graphs $G \in \mathbb{G}_V$ to an ordering $\preceq_G^F \in L(V)$. If F is a partial function then it is called a *partial ranking system*, otherwise it is called a *general ranking system*.

One can view this setting as a variation/extension of the classical theory of social choice as modeled by [Arrow, 1963]. The ranking systems setting differs in two main properties. First, in this setting we assume that the set of voters and the set of alternatives coincide, and second, we allow agents only two levels of preference over the alternatives, as opposed to Arrow’s setting where agents could rank alternatives arbitrarily.

3 Transitivity

A basic property one would assume of ranking systems is that if an agent a ’s voters are ranked higher than those of agent b ,




Figure 1: Example of Transitivity

then agent a should be ranked higher than agent b . This notion is formally captured below:

Definition 3.1. Let F be a ranking system. We say that F satisfies *strong transitivity* if for all graphs $G = (V, E)$ and for all vertices $v_1, v_2 \in V$: Assume there is a 1-1 mapping $f : P(v_1) \mapsto P(v_2)$ s.t. for all $v \in P(v_1)$: $v \preceq f(v)$. Further assume that either f is not onto or for some $v \in P(v_1)$: $v \prec f(v)$. Then, $v_1 \prec v_2$.

Consider for example the graph G in Figure 1 and any ranking system F that satisfies strong transitivity. F must rank vertex d below all other vertices, as it has no predecessors, unlike all other vertices. If we assume that $a \preceq_G^F b$, then by strong transitivity we must conclude that $b \preceq_G^F c$ as well. But then we must conclude that $b \prec_G^F a$ (as b ’s predecessor a is ranked lower than a ’s predecessor c , and a has an additional predecessor d), which leads to a contradiction. Given $b \prec_G^F a$, again by transitivity, we must conclude that $c \prec_G^F b$, so the only ranking for the graph G that satisfies strong transitivity is $d \prec_G^F c \prec_G^F b \prec_G^F a$.

In [Tennenholtz, 2004], we have suggested an algorithm that defines a ranking system that satisfies strong transitivity by iteratively refining an ordering of the vertices.

Note that the PageRank ranking system defined in [Page et al., 1998] does not satisfy strong transitivity. This is due to the fact that PageRank reduces the weight of links (or votes) from nodes which have a higher out-degree. Thus, assuming Yahoo! and Microsoft are equally ranked, a link from Yahoo! means less than a link from Microsoft, because Yahoo! links to more external pages than does Microsoft. Noting this fact, we can weaken the definition of transitivity to require that the predecessors of the compared agents have an equal out-degree:

Definition 3.2. Let F be a ranking system. We say that F satisfies *weak transitivity* if for all graphs $G = (V, E)$ and for all vertices $v_1, v_2 \in V$: Assume there is a 1-1 mapping $f : P(v_1) \mapsto P(v_2)$ s.t. for all $v \in P(v_1)$: $v \preceq f(v)$ and $|S(v)| = |S(f(v))|$. Further assume that either f is not onto or for some $v \in P(v_1)$: $v \prec f(v)$. Then, $v_1 \prec v_2$.

Indeed, an idealized version of the PageRank ranking system defined on strongly connected graphs satisfies this weakened version of transitivity. Furthermore, the result in the example above does not change when we consider weak transitivity in place of strong transitivity.

4 Ranked Independence of Irrelevant Alternatives

A standard assumption in social choice settings is that an agent’s relative rank should only depend on (some property of) their immediate predecessors. Such axioms are usually called independence of irrelevant alternatives (IIA) axioms.




Figure 2: An example of RIIA.

In our setting, we require the relative ranking of two agents must only depend on the pairwise comparisons of the ranks of their predecessors, and not on their identity or cardinal value. Our IIA axiom, called *ranked IIA*, differs from the one suggested by [Arrow, 1963] in the fact that we do not consider the identity of the voters, but rather their relative rank.

For example, consider the graph in Figure 2. Furthermore, assume a ranking system F has ranked the vertices of this graph as following: $a \simeq b \prec c \simeq d \prec e \simeq f$. Now look at the comparison between c and d . c 's predecessors, a and b , are both ranked equally, and both ranked lower than d 's predecessor f . This is also true when considering e and f – e 's predecessors c and d are both ranked equally, and both ranked lower than f 's predecessor e . Therefore, if we agree with ranked IIA, the relation between c and d , and the relation between e and f must be the same, which indeed it is – both $c \simeq d$ and $e \simeq f$. However, this same situation also occurs when comparing c and f (c 's predecessors a and b are equally ranked and ranked lower than f 's predecessor e), but in this case $c \prec f$. So, we can conclude that the ranking system F which produced these rankings does not satisfy ranked IIA.

To formally define this condition, one must consider all possibilities of comparing two nodes in a graph based only on ordinal comparisons of their predecessors. We call these possibilities comparison profiles:

Definition 4.1. A *comparison profile* is a pair $\langle \mathbf{a}, \mathbf{b} \rangle$ where $\mathbf{a} = (a_1, \dots, a_n)$, $\mathbf{b} = (b_1, \dots, b_m)$, $a_1, \dots, a_n, b_1, \dots, b_m \in \mathbb{N}$, $a_1 \leq a_2 \leq \dots \leq a_n$, and $b_1 \leq b_2 \leq \dots \leq b_m$. Let \mathcal{P} be the set of all such profiles.

A ranking system F , a graph $G = (V, E)$, and a pair of vertices $v_1, v_2 \in V$ are said to *satisfy* such a comparison profile $\langle \mathbf{a}, \mathbf{b} \rangle$ if there exist 1-1 mappings $f_1 : P(v_1) \mapsto \{1 \dots n\}$ and $f_2 : P(v_2) \mapsto \{1 \dots m\}$ such that given $f : (\{1\} \times P(v_1)) \cup (\{2\} \times P(v_2)) \mapsto \mathbb{N}$ defined as:

$$\begin{aligned} f(1, v) &= a_{f_1(v)} \\ f(2, u) &= b_{f_2(u)}, \end{aligned}$$

$$f(i, x) \leq f(j, y) \Leftrightarrow x \preceq_G^F y \text{ for all } (i, x), (j, y) \in (\{1\} \times P(v_1)) \cup (\{2\} \times P(v_2)).$$

For example, in the example considered above, all of the pairs (c, d) , (c, f) , and (e, f) satisfy the comparison profile $\langle (1, 1), (2) \rangle$.

We now require that for every such profile the ranking system ranks the nodes consistently:

Definition 4.2. Let F be a ranking system. We say that F satisfies *ranked independence of irrelevant alternatives (RIIA)* if there exists a mapping $f : \mathcal{P} \mapsto \{0, 1\}$ such that for every graph $G = (V, E)$ and for every pair of vertices $v_1, v_2 \in V$

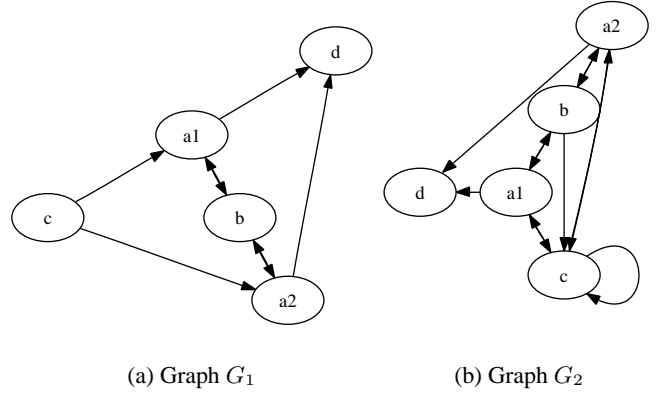


Figure 3: Graphs for the proof of Theorem 5.1

and for every comparison profile $p \in \mathcal{P}$ that v_1 and v_2 satisfy, $v_1 \preceq_G^F v_2 \Leftrightarrow f(p) = 1$.

As RIIA is an independence property, the ranking system F_{\equiv} , that ranks all agents equally, satisfies RIIA. A more interesting ranking system that satisfies RIIA is the approval voting ranking system, defined below.

Definition 4.3. The *approval voting ranking system AV* is the ranking system defined by:

$$v_1 \preceq_G^{AV} v_2 \Leftrightarrow |P(v_1)| \leq |P(v_2)|$$

A full axiomatization of the approval voting ranking system is given in section 7.

5 Impossibility

Our main result illustrates the impossibility of satisfying (weak) transitivity and RIIA simultaneously.

Theorem 5.1. *There is no general ranking system that satisfies weak transitivity and RIIA.*

Proof. Assume for contradiction that there exists a ranking system F that satisfies weak transitivity and RIIA. Consider first the graph G_1 in Figure 3(a). First, note that a_1 and a_2 satisfy some comparison profile $p_a = ((x, y), (x, y))$ because they have identical predecessors. Thus, by RIIA, $a_1 \preceq_{G_1}^F a_2 \Leftrightarrow a_2 \preceq_{G_1}^F a_1$, and therefore $a_1 \simeq_{G_1}^F a_2$. By weak transitivity, it is easy to see that $c \prec_{G_1}^F a_1$ and $c \prec_{G_1}^F b$. If we assume $b \preceq_{G_1}^F a_1$, then by weak transitivity, $a_1 \prec_{G_1}^F b$ which contradicts our assumption. So we conclude that $c \prec_{G_1}^F a_1 \prec_{G_1}^F b$.

Now consider the graph G_2 in Figure 3(b). Again, by RIIA, $a_1 \simeq_{G_2}^F a_2$. By weak transitivity, it is easy to see that $a_1 \prec_{G_2}^F c$ and $b \prec_{G_2}^F c$. If we assume $a_1 \preceq_{G_2}^F b$, then by weak transitivity, $b \prec_{G_2}^F a_1$ which contradicts our assumption. So we conclude that $b \prec_{G_2}^F a_1 \prec_{G_2}^F c$.

Consider the comparison profile $p = ((1, 3), (2, 2))$. Given F , a_1 and b satisfy p in G_1 (because $c \prec_{G_1}^F a_1 \simeq_{G_1}^F a_2 \prec_{G_1}^F a_1$) and in G_2 (because $b \prec_{G_2}^F a_1 \simeq_{G_2}^F a_2 \prec_{G_2}^F c$). Thus, by

RIIA, $a_1 \preceq_{G_1}^F b \Leftrightarrow a_1 \preceq_{G_2}^F b$, which is a contradiction to the fact that $a_1 \prec_{G_1}^F b$ but $b \prec_{G_2}^F a_1$. \square

This result is quite a surprise, as it means that every reasonable definition of a ranking system must either consider cardinal values for nodes and/or edges (like [Page *et al.*, 1998]), or operate ordinally on a global scale (like [Altman and Tennenholtz, 2005]).

6 Relaxing Generality

A hidden assumption in our impossibility result is the fact that we considered only general ranking systems. In this section we analyze several special classes of graphs that relate to common ranking scenarios.

6.1 Small Graphs

A natural limitation on a preference graph is a cap on the number of vertices (agents) that participate in the ranking. Indeed, when there are three or less agents involved in the ranking, strong transitivity and RIIA can be simultaneously satisfied. An appropriate ranking algorithm for this case is the one we suggested in [Tennenholtz, 2004].

However, when there are four or more agents, strong transitivity and RIIA cannot be simultaneously satisfied (the proof is similar to that of Theorem 5.1, but with vertex d removed in both graphs). When five or more agents are involved, even weak transitivity and RIIA cannot be simultaneously satisfied, as implied by the proof of Theorem 5.1.

6.2 Single Vote Setting

Another natural limitation on the domain of graphs that we might be interested in is the restriction of each agent(vertex) to exactly one vote(successor). For example, in the voting paradigm this could be viewed as a setting where every agent votes for exactly one agent. The following proposition shows that even in this simple setting weak transitivity and RIIA cannot be simultaneously satisfied.

Proposition 6.1. *Let \mathbb{G}_1 be the set of all graphs $G = (V, E)$ such that $|S(v)| = 1$ for all $v \in V$. There is no partial ranking system over \mathbb{G}_1 that satisfies weak transitivity and RIIA.*

Proof. Assume for contradiction that there is a partial ranking system F over \mathbb{G}_1 that satisfies weak transitivity and RIIA. Let $f : \mathcal{P} \mapsto \{0, 1\}$ be the mapping from the definition of RIIA for F .

Let $G_1 \in \mathbb{G}_1$ be the graph in Figure 4a. By weak transitivity, $x_1 \simeq_{G_1}^F x_2 \prec_{G_1}^F b \prec_{G_1}^F a$. (a, b) satisfies the comparison profile $\langle (1, 1, 2), (3) \rangle$, so we must have $f \langle (1, 1, 2), (3) \rangle = 0$. Now let $G_2 \in \mathbb{G}_1$ be the graph in Figure 4b. By weak transitivity $x_1 \simeq_{G_2}^F x_2 \prec_{G_2}^F y \prec_{G_2}^F a \prec_{G_2}^F b$. (b, a) satisfies the comparison profile $\langle (2, 3), (1, 4) \rangle$, so we must have $f \langle (2, 3), (1, 4) \rangle = 0$.

Let $G_3 \in \mathbb{G}_1$ be the graph in Figure 4c. By weak transitivity it is easy to see that $x_1 \simeq_{G_3}^F \dots \simeq_{G_3}^F x_7 \prec_{G_3}^F y_1 \simeq_{G_3}^F y_2 \prec_{G_3}^F c \prec_{G_3}^F d$. Furthermore, by weak transitivity we conclude that $a \prec_{G_3}^F b$ and $a' \prec_{G_3}^F b'$ from $c \prec_{G_3}^F d$; and

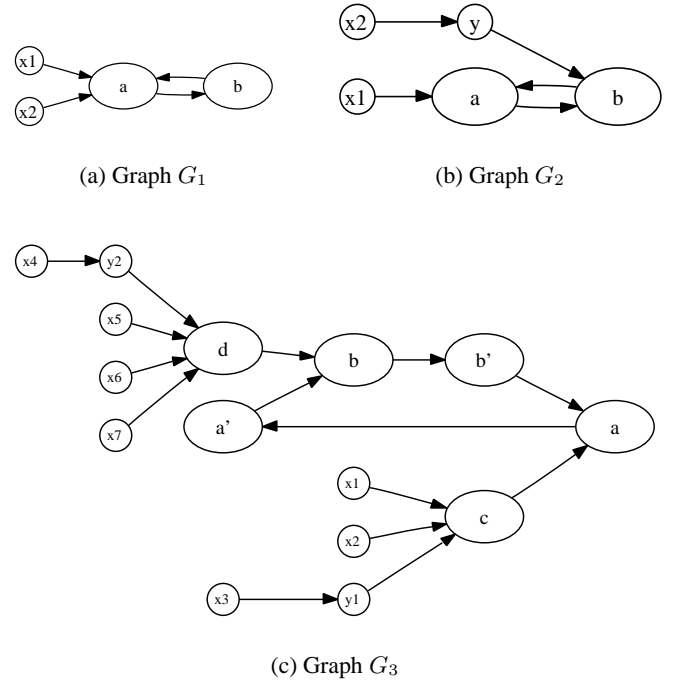


Figure 4: Graphs from the proof of proposition 6.1

$y_1 \prec_{G_3}^F b$ from $x_3 \prec_{G_3}^F d$. Now consider the vertex pair (c, b') . We have shown that $x_1 \simeq_{G_3}^F x_2 \prec_{G_3}^F y_1 \prec_{G_3}^F b$. So, (c, b') satisfies the comparison profile $\langle (1, 1, 2), (3) \rangle$, thus by RIIA $b' \prec_{G_3}^F c$. Now consider the vertex pair (b, a) . We have already shown that $a' \prec_{G_3}^F b' \prec_{G_3}^F c \prec_{G_3}^F d$. So, (a, b) satisfies the comparison profile $\langle (2, 3), (1, 4) \rangle$, thus by RIIA $b \prec_{G_3}^F a$. However, we have already shown that $a \prec_{G_3}^F b$ – a contradiction. Thus, the ranking system F cannot exist. \square

6.3 Bipartite Setting

In the world of reputation systems [Resnick *et al.*, 2000], we frequently observe a distinction between two types of agents such that each type of agent only ranks agents of the other type. For example buyers only interact with sellers and vice versa. This type of limitation is captured by requiring the preference graphs to be bipartite, as defined below.

Definition 6.2. A graph $G = (V, E)$ is called *bipartite* if there exist V_1, V_2 such that $V = V_1 \cup V_2$, $V_1 \cap V_2 = \emptyset$, and $E \subseteq (V_1 \times V_2) \cup (V_2 \times V_1)$. Let \mathbb{G}_B be the set of all bipartite graphs.

Our impossibility result extends to the limited domain of bipartite graphs.

Proposition 6.3. *There is no partial ranking system over $\mathbb{G}_B \cap \mathbb{G}_1$ that satisfies weak transitivity and RIIA.*

6.4 Strongly Connected Graphs

The well-known PageRank ranking system is (ideally) defined on the set of strongly connected graphs. That is, the

set of graphs where there exists a directed path between any two vertices.

Let us denote the set of all strongly connected graphs \mathbb{G}_{SC} . The following proposition extends our impossibility result to strongly connected graphs.

Proposition 6.4. *There is no partial ranking system over \mathbb{G}_{SC} .*

7 Axiomatization of Approval Voting

In the previous sections we have seen mostly negative results which arise when trying to accommodate (weak) transitivity and RIIA. We have shown that although each of the axioms can be satisfied separately, there exists no general ranking system that satisfies both axioms.

We have previously shown [Tennenholtz, 2004] a non-trivial ranking system that satisfies (weak) transitivity, but have not yet shown such a ranking system for RIIA. In this section we provide a representation theorem for a ranking system that satisfies RIIA but not weak transitivity — the approval voting ranking system. This system ranks the agents based on the number of votes each agent received, with no regard to the rank of the voters. The axiomatization we provide in this section shows the power of RIIA, as it shows that there exists only one (interesting) ranking system that satisfies it without introducing transitive effects.

In order to specify our axiomatization, recall the following classical definitions from the theory of social choice:

The positive response axiom essentially means that if an agent receives additional votes, its rank must improve:

Definition 7.1. Let F be a ranking system. F satisfies *positive response* if for all graphs $G = (V, E)$ and for all $(v_1, v_2) \in (V \times V) \setminus E$, and for all $v_3 \in V$: Let $G' = (V, E \cup (v_1, v_2))$. If $v_3 \preceq_G^F v_2$, then $v_3 \prec_{G'}^F v_2$.

The anonymity and neutrality axioms mean that the names of the voters and alternatives respectively do not matter for the ranking:

Definition 7.2. A ranking system F satisfies *anonymity* if for all $G = (V, E)$, for all permutations $\pi : V \mapsto V$, and for all $v_1, v_2 \in V$: Let $E' = \{(\pi(v_1), v_2) | (v_1, v_2) \in E\}$. Then, $v_1 \preceq_{(V, E)}^F v_2 \Leftrightarrow v_1 \preceq_{(V, E')}^F v_2$.

Definition 7.3. A ranking system F satisfies *neutrality* if for all $G = (V, E)$, for all permutations $\pi : V \mapsto V$, and for all $v_1, v_2 \in V$: Let $E' = \{(v_1, \pi(v_2)) | (v_1, v_2) \in E\}$. Then, $v_1 \preceq_{(V, E)}^F v_2 \Leftrightarrow v_1 \preceq_{(V, E')}^F v_2$.

Arrow's classical Independence of Irrelevant Alternatives axiom requires that the relative rank of two agents be dependent only on the set of agents that preferred one over the other.

Definition 7.4. A ranking system F satisfies *Arrow's Independence of Irrelevant Alternatives (AIIA)* if for all $G = (V, E)$, for all $G' = (V, E')$, and for all $v_1, v_2 \in V$: Let $P_G(v_1) \setminus P_G(v_2) = P_{G'}(v_1) \setminus P_{G'}(v_2)$ and $P_G(v_2) \setminus P_G(v_1) = P_{G'}(v_2) \setminus P_{G'}(v_1)$. Then, $v_1 \preceq_G^F v_2 \Leftrightarrow v_1 \preceq_{G'}^F v_2$.

Our representation theorem states that together with positive response and RIIA, any one of the three independence

conditions above (anonymity, neutrality, and AIIA) are essential and sufficient for a ranking system being AV^1 . In addition, we show that as in the classical social choice setting when only considering two-level preferences, positive response, anonymity, neutrality, and AIIA are an essential and sufficient representation of approval voting. This result extends the well known axiomatization of the majority rule due to [May, 1952]:

Proposition 7.5. (*May's Theorem*) *A social welfare functional over two alternatives is a majority social welfare functional if and only if it satisfies anonymity, neutrality, and positive response.*

We can now formally state our theorem:

Theorem 7.6. *Let F be a general ranking system. Then, the following statements are equivalent:*

1. F is the approval voting ranking system ($F = AV$)
2. F satisfies positive response, anonymity, neutrality, and AIIA
3. F satisfies positive response, RIIA, and either one of anonymity, neutrality, and AIIA

Proof. (Sketch) It is easy to see that AV satisfies positive response, RIIA, anonymity, neutrality, and AIIA. It remains to show that (2) and (3) entail (1) above.

To prove (2) entails (1), assume that F satisfies positive response, anonymity, neutrality, and AIIA. Let $G = (V, E)$ be some graph and let $v_1, v_2 \in V$ be some agents. By AIIA, the relative ranking of v_1 and v_2 depends only on the sets $P_G(v_1) \setminus P_G(v_2)$ and $P_G(v_2) \setminus P_G(v_1)$. We have now narrowed our consideration to a set of agents with preferences over two alternatives, so we can apply Proposition 7.5 to complete our proof.

To prove (3) entails (1), assume that F satisfies positive response, RIIA and either anonymity or neutrality or AIIA. As F satisfies RIIA we can limit our discussion to comparison profiles. Let $f : \mathcal{P} \mapsto \{0, 1\}$ be the function from the definition of RIIA. We will use the notation $\mathbf{a} \preceq \mathbf{b}$ to mean $f\langle \mathbf{a}, \mathbf{b} \rangle = 1$, $\mathbf{a} \prec \mathbf{b}$ to mean $f\langle \mathbf{b}, \mathbf{a} \rangle = 0$, and $\mathbf{a} \simeq \mathbf{b}$ to mean $\mathbf{a} \preceq \mathbf{b}$ and $\mathbf{b} \preceq \mathbf{a}$.

By the definition of RIIA, it is easy to see that $\mathbf{a} \simeq \mathbf{a}$ for all \mathbf{a} . By positive response it is also easy to see that $\underbrace{(1, 1, \dots, 1)}_n \preceq \underbrace{(1, 1, \dots, 1)}_m$ iff $n \leq m$. Let $P = \langle \langle a_1, \dots, a_n \rangle, \langle b_1, \dots, b_m \rangle \rangle$ be a comparison profile. Let $G = (V, E)$ be the following graph :

$$\begin{aligned} V &= \{x_1, \dots, x_{\max\{a_n, b_m\}}\} \cup \\ &\quad \cup \{v_1, \dots, v_n, v'_1, \dots, v'_n, v\} \cup \\ &\quad \cup \{u_1, \dots, u_m, u'_1, \dots, u'_m, u\} \\ E &= \{(x_i, v_j) | i \leq a_j\} \cup \{(x_i, u_j) | i \leq b_j\} \cup \\ &\quad \cup \{(v_i, v) | i = 1, \dots, n\} \cup \{(u_i, u) | i = 1, \dots, m\}. \end{aligned}$$

¹In fact, an even weaker condition of *decoupling*, that in essence allows us to permute the graph structure while keeping the edges' names is sufficient in this case.

It is easy to see that in the graph G , v and u satisfy the profile P . Let π be the following permutation:

$$\pi(x) = \begin{cases} v'_i & x = v_i \\ v_i & x = v'_i \\ u'_i & x = u_i \\ u_i & x = u'_i \\ x & \text{Otherwise.} \end{cases}$$

The remainder of the proof depends on which additional axiom F satisfies:

- If F satisfies anonymity, let $E' = \{(\pi(x), y) | (x, y) \in E\}$. Note that in the graph (V, E') v and u satisfy the profile $\langle \underbrace{(1, 1, \dots, 1)}_n, \underbrace{(1, 1, \dots, 1)}_m \rangle$, and thus $v \preceq_{(V, E')}^F u$.
 $u \Leftrightarrow n \leq m$. By anonymity, $u \preceq_{(V, E)}^F v \Leftrightarrow u \preceq_{(V, E')}^F v$, thus proving that $f(P) = 1 \Leftrightarrow n \leq m$ for an arbitrary comparison profile P , and thus $F = AV$.
- If F satisfies neutrality, let $E' = \{(x, \pi(y)) | (x, y) \in E\}$. Note that in the graph (V, E') v and u satisfy the profile $\langle \underbrace{(1, 1, \dots, 1)}_n, \underbrace{(1, 1, \dots, 1)}_m \rangle$, and thus $v \preceq_{(V, E')}^F u$.
 $u \Leftrightarrow n \leq m$. By neutrality, $u \preceq_{(V, E)}^F v \Leftrightarrow u \preceq_{(V, E')}^F v$, again showing that $f(P) = 1 \Leftrightarrow n \leq m$ for an arbitrary comparison profile P , and thus $F = AV$.
- If F satisfies AIIA, let $E' = \{(x, \pi(y)) | (x, y) \in E\}$ as before. So, also $v \preceq_{(V, E')}^F u \Leftrightarrow n \leq m$. Note that $P_G(v) = P_{(V, E')}(v)$ and $P_G(u) = P_{(V, E')}(u)$, so by AIIA, $u \preceq_{(V, E)}^F v \Leftrightarrow u \preceq_{(V, E')}^F v$, and thus as before, $F = AV$. \square

8 Concluding Remarks

Reasoning about preferences and preference aggregation is a fundamental task in reasoning about multi-agent systems (see e.g. [Boutilier *et al.*, 2004; Conitzer and Sandholm, 2002; LaMura and Shoham, 1998]). A typical instance of preference aggregation is the setting of ranking systems. Ranking systems are fundamental ingredients of some of the most famous tools/techniques in the Internet (e.g. Google's page rank and eBay's reputation systems, among many others).

Our aim in this paper was to treat ranking systems from an axiomatic perspective. The classical theory of social choice lay the foundations to a large part of the rigorous work on multi-agent systems. Indeed, the most classical results in the theory of mechanism design, such as the Gibbard-Satterthwaite Theorem (see [Mas-Colell *et al.*, 1995]) are applications of the theory of social choice. Moreover, previous work in AI has employed the theory of social choice for obtaining foundations for reasoning tasks [Doyle and Wellman, 1989] and multi-agent coordination [Kfir-Dahav and Tennenholtz, 1996]. It is however interesting to note that ranking systems suggest a novel and new type of theory of social choice. We see this point as especially attractive, and as a main reason for concentrating on the study of the axiomatic foundations of ranking systems.

In this paper we identified two fundamental axioms for ranking systems, and conducted a basic axiomatic study of

such systems. In particular, we presented surprising impossibility results, and a representation theorem for the well-known approval voting scheme.

References

- [Altman and Tennenholtz, 2005] Alon Altman and Moshe Tennenholtz. Ranking systems: The pagerank axioms. In *Proceedings of the 6th ACM conference on Electronic commerce (EC-05)*, 2005.
- [Arrow, 1963] K.J. Arrow. *Social Choice and Individual Values (2nd Ed.)*. Yale University Press, 1963.
- [Boutilier *et al.*, 2004] Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21:135–191, 2004.
- [Conitzer and Sandholm, 2002] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the 18th conference on uncertainty in Artificial Intelligence (UAI-02)*, pages 103–110, 2002.
- [Doyle and Wellman, 1989] J. Doyle and M.P. Wellman. Impediments to Universal Preference-Based Default Theories. In *Proceedings of the 1st conference on principles of knowledge representation and reasoning*, 1989.
- [Kfir-Dahav and Tennenholtz, 1996] Noa E. Kfir-Dahav and Moshe Tennenholtz. Multi-Agent Belief Revision. In *Proceedings of the 6th conference on theoretical aspects of rationality and knowledge (TARK)*, 1996.
- [LaMura and Shoham, 1998] P. LaMura and Y. Shoham. Conditional, Hierarchical Multi-Agent Preferences. In *Proceedings of Theoretical Aspects of Rationality and Knowledge*, pages 215–224, 1998.
- [Mas-Colell *et al.*, 1995] A. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [May, 1952] Kenneth O. May. A set of independent, necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4):680–84, 1952.
- [Page *et al.*, 1998] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report, Stanford University, 1998.
- [Resnick and Zeckhauser, 2001] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. Working Paper for the NBER workshop on empirical studies of electronic commerce, January 2001.
- [Resnick *et al.*, 2000] P. Resnick, R. Zeckhauser, R. Friedman, and E. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.
- [Tennenholtz, 2004] M. Tennenholtz. Reputation systems: An axiomatic approach. In *Proceedings of the 20th conference on uncertainty in Artificial Intelligence (UAI-04)*, 2004.