Journal of Artificial Intelligence Research 31 (2008) 473-495 Submitted 03/07; published 03/08

Axiomatic Foundations for Ranking Systems

Alon Altman EPSALON@STANFORD.EDU
Department of Computer Science
Stanford University
Stanford, CA 94305-9020 USA

Moshe Tennenholtz MOSHET@IE.TECHNION.AC.IL
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. However, when transitivity is weakened, an interesting possibility result is obtained. In addition, we show a complete axiomatization of approval voting using ranked IIA.

1. Introduction

The ranking of agents based on other agents' input is fundamental to multi-agent systems (see e.g. Resnick, Zeckhauser, Friedman, & Kuwabara, 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, Brin, Motwani, & Winograd, 1998) and eBay's reputation system (Resnick & 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.

Much effort has 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; see Moulin, 1991 for an overview). In the ranking systems setting, we have successfully applied the descriptive perspective by providing a representation theorem (Altman & Tennenholtz, 2005b) for the well-known PageRank algorithm (Page et al., 1998), which is the basis of Google's search technology (Brin & Page, 1998).

An excellent example for the normative perspective is Arrow's impossibility theorem mentioned above. Tennenholtz (2004) has 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. Still in the normative approach to ranking systems, we have tackled the issue of incentives (Altman & Tennenholtz, 2006b, 2006c), with both positive and negative results. Recently, we have considered a variation of ranking systems, where a personalized ranking is generated for every participant in the system (Altman & Tennenholtz, 2006a), with surprisingly different results.

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! This result means that if we would like to fully capture transitive effects, ranking decisions must be made globally, or be based on numeric calculations. We further show that our impossibility result holds under various restrictions on the class of ranking problems considered.

On the other hand, we show a positive result for the case when the transitivity axiom is relaxed. This new ranking system is practical and useful and an algorithm is provided for its computation. Finally, 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 main positive result, in the form of a ranking system satisfying a weaker version of transitivity is given in Section 7, while an axiomatization for the Approval Voting ranking system is presented in Section 8. Finally, some concluding remarks are given in Section 9.

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 directed 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.

The two-level limitation is important in order to avoid Arrow-style impossibility results. Indeed, in the dichotomous (i.e. two level) setting such results do not apply (Bogomolnaia, Moulin, & Stong, 2005). Had we allowed general rankings as the input of the system, we would have reached impossibility results as a direct result of Arrow-style impossibility. By adding the dichotomous limitation, we ensure that our results will be a consequence of the co-incidence of the voters and alternatives and the related transitive effects.

2.1 Examples of Ranking Systems

In order to make the abstract definition of ranking systems above more concrete, we shall now give some examples of several well-known ranking systems. In order to define these systems, and throughout this paper, we shall use the following notation:

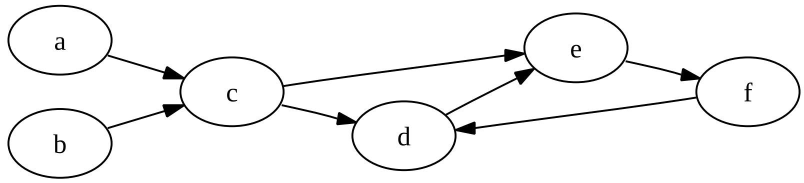

Figure 1: Example graph for ranking systems.

Notation 2.4. Let $G = (V, E)$ be some graph and $v \in V$ be some vertex. Let $P_G(v) \triangleq \{u|(u, v) \in E\}$ and $S_G(v) \triangleq \{u|(v, u) \in E\}$ denote the predecessor and successor sets of v in G respectively. When G is understood from context, we will sloppily use $P(v)$ and $S(v)$.

Approval Voting is a very simple ranking system that ranks the agents according to the number of votes (i.e. incoming edges) they have. Formally,

Definition 2.5. The *approval voting ranking system AV* is the ranking system defined by:

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

Consider the graph in Figure 1. The AV ranking system would rank this graph $a \simeq b \prec f \prec c \simeq d \simeq e$ based on the fact that the vertices in $\{a, b\}$, $\{f\}$, and $\{c, d, e\}$ have 0, 1, and 2 predecessors respectively. A full axiomatization of the approval voting ranking system is given in section 8.

One major application of Ranking Systems is in the context of Internet pages. In that context, we represent the Internet as a directed graph, where the vertices are websites, and the edges are links between these websites. A prominent ranking system in this setting is PageRank (Page et al., 1998), which is based on a random walk of the Internet graph. Namely, in this process we start in a random page, and iteratively move to one of the pages that are linked to by the current page, assigning equal probabilities to each such page. We define the PageRank matrix which captures the random walk created by the PageRank procedure:

Definition 2.6. Let $G = (V, E)$ be a graph, and assume $V = \{v_1, v_2, \ldots, v_n\}$. The *PageRank Matrix* A_G (of dimension $n \times n$) is defined as:

$$[A_G]_{i,j} = \begin{cases} 1/|S_G(v_j)| & (v_j, v_i) \in E \\ 0 & \text{Otherwise.} \end{cases}$$

The PageRank procedure will rank pages according to the stationary probability distribution obtained in the limit of the above random walk; this is formally defined as follows:

Definition 2.7. Let $G = (V, E)$ be some graph, and assume $V = \{v_1, v_2, \ldots, v_n\}$. Let $0 \leq d < 1$ be a *damping factor*. Let $\mathbf{r}$ be the unique solution of the system $(1-d) \cdot A_G \cdot \mathbf{r} + d \cdot \begin{pmatrix} 1 & 1 & \cdots & 1 \end{pmatrix}^T = \mathbf{r}$ where $\sum r_i = n$. If there is no unique solution, then the ranking is not defined. Otherwise, the *PageRank* $PR_G(v_i)$ of a vertex $v_i \in V$ is defined as $PR_G(v_i) = r_i$. The *PageRank ranking system* is a ranking system that for the vertex set V maps G to $\preceq_G^{PR}$, where $\preceq_G^{PR}$ is defined as: for all $v_i, v_j \in V$: $v_i \preceq_G^{PR} v_j$ if and only if $PR_G(v_i) \leq PR_G(v_j)$.

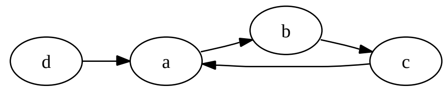

Figure 2: Example of Transitivity

It can be shown that for $d > 0$, there is indeed a unique solution and thus the ranking system is a general one. However, when $d = 0$ this ranking system becomes a partial ranking system, as it is not always well defined.

In the graph in Figure 1, for $d = 0.2$ the PageRank values assigned for $a \ldots f$ are $(0.2, 0.2, 0.52, 1.7, 1.77, 1.61)$ giving the ranking $a \simeq b \prec c \prec f \prec d \prec e$. Note that this ranking differs from the one assigned by approval voting, that neither of the rankings is a refinement of the other. This example shows that PageRank and Approval Voting are distinct ranking systems, and that the two may disagree on the ranking of two vertices. We will soon see that these systems satisfy two mutually exclusive properties of ranking systems.

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, 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 (but not necessarily onto) $f : P(v_1) \mapsto P(v_2)$ s.t. for all $v \in P(v_1)$: $v \preceq f(v)$. Then, $v_1 \preceq v_2$. Further assume that either f is not onto or for some $v \in P(v_1)$: $v \prec f(v)$. Then, $v_1 \prec v_2$.

To explain how the formal definition captures the intuition, consider the simple graph

$$a \rightarrow b \rightarrow c$$

Our intuition tells us that c who is at the end of the vote chain should be more trusted, and thus ranked higher than b, this is because of the fact that b is more trusted than a, due to b having a vote compared to a having none. This intuition is correctly captured by the definition above: a must be ranked strictly below b because any function mapping $P(a) = \emptyset$ to $P(b) = \{a\}$ is not onto, and b must be ranked strictly below c because the trivial mapping from $P(b) = \{a\}$ to $P(c) = \{b\}$ satisfies $a \prec b$, and thus we get $b \prec c$, as expected.

For a more involved example, consider the graph G in Figure 2 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 $a \preceq^F_G b$, then by strong transitivity we must conclude that $b \preceq^F_G c$ as well. But then we must conclude that $b \prec^F_G a$ (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 $b \prec^F_G a$, again by transitivity, we must conclude that $c \prec^F_G b$, so the only ranking for the graph G that satisfies strong transitivity is $d \prec^F_G c \prec^F_G b \prec^F_G a$.

Tennenholtz (2004) has suggested an algorithm that defines a ranking system that satisfies strong transitivity by iteratively refining an ordering of the vertices starting from the ranking suggested by approval voting.

Note that the PageRank ranking system 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 $|S(v)| = |S(f(v))|$. Then, $v_1 \preceq v_2$. Further assume that either f is not onto or for some $v \in P(v_1)$: $v \prec f(v)$. Then, $v_1 \prec v_2$.

For an example of weak transitivity, one can reconsider the strong transitivity example above, as it still applies to weak transitivity.

The PageRank ranking system satisfies this weakened version of transitivity. This is due to the fact that:

$$PR(v_1) = \sum_{v \in P(v_1)} \frac{PR(v)}{|S(v)|} \leq \sum_{v \in P(v_1)} \frac{PR(f(v))}{|S(f(v))|} \leq \sum_{v \in P(v_2)} \frac{PR(v)}{|S(v)|} = PR(v_2).$$

In the case where $v \prec f(v)$ for some $v \in P(v_1)$ the first inequality is strict, and if f is not onto the second inequality is strict.

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) the agents who have voted for them. Such axioms are usually called independence of irrelevant alternatives (IIA) axioms. In our setting, such IIA axioms mean that an agent's rank must only depend on a property of its immediate predecessors.

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 3. 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 $c \prec f$. All three cases involve comparing two vertices, one with two weaker predecessors and one with one stronger

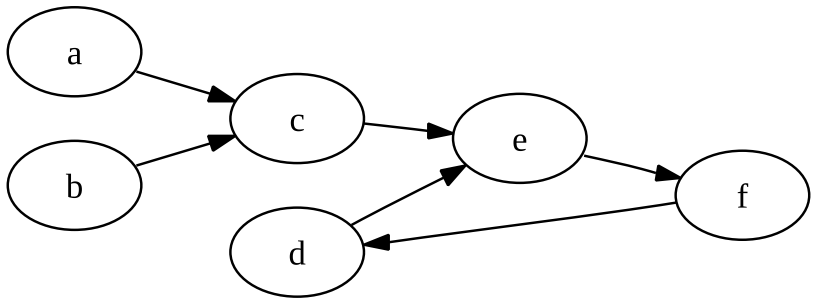

Figure 3: An example of RIIA.

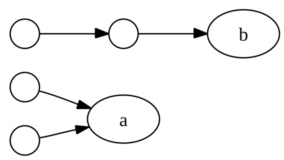

Figure 4: Graph for profile $\langle (1,1),(2) \rangle$.

predecessor, but the outcome of these comparisons in not consistent. Therefore, 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 $\mathbf{a} = (a_1, \ldots, a_n)$, $\mathbf{b} = (b_1, \ldots, b_m)$, $a_1, \ldots, a_n, b_1, \ldots, b_m \in \mathbb{N}$, $a_1 \le a_2 \le \cdots \le a_n$, and $b_1 \le b_2 \le \cdots \le b_m$. 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 $f_1 : P(v_1) \mapsto \{1 \ldots n\}$ and $f_2 : P(v_2) \mapsto \{1 \ldots m\}$ such that given $f : (\{1\} \times P(v_1)) \cup (\{2\} \times P(v_2)) \mapsto \mathbb{N}$ defined as:

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

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

Consider the profile $\langle (1,1),(2) \rangle$. This comparison profile illustrates the basic question of comparing an agent who got two low-rank votes with one who got one high-rank vote. This question is undecided by transitivity alone, and if we do assume transitivity this comparison profile is satisfied by the pair (a, b) in the graph in Figure 4. The f function above simply maps the predecessors of a and b to 1 and 2 respectively.

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

Notation 4.3. When the function f from the definition above is understood from context, 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 $\mathbf{b} \preceq \mathbf{a}$.

For example, in the example considered above, all of the pairs (c, d), (c, f), and (e, f) satisfy the comparison profile $\langle (1, 1), (2) \rangle$. As we have seen above, the pairs (c, d) and (e, f) entail that $(1, 1) \simeq (2)$, while (c, f) entails that $(1, 1) \prec (2)$. These results contradict each other, and therefore we conclude that the ranking system that produced this ranking does not satisfy RIIA.

The definition of RIIA formalizes the requirement of consistency in the comparisons such as the one we have seen above. It means that any ranking system satisfying RIIA must decide on the relative ranking of a and b in Figure 4, and (assuming transitivity) rank the same in all other occurrences of two weak vs. one strong predecessor.

As RIIA is an independence property, the ranking system $F_=$, that ranks all agents equally, satisfies RIIA.

The approval voting ranking system AV also satisfies RIIA. This is due to the fact that for any comparison profile $\langle (a_1, \ldots, a_n), (b_1, \ldots b_m) \rangle$, the f function for AV ranks $\mathbf{a} \preceq \mathbf{b}$ iff $n \leq m$. We will use this fact in the axiomatization of approval voting we present in Section 8.

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 5(a). Note that all vertices in this graph have an out-degree of 2 or 0, and thus the out-degree requirement of weak transitivity is trivially fulfilled. Now note that a_1 and a_2 must satisfy some comparison profile $p_a = ((x, y), (x, y))$ because they have identical predecessors. Thus, by RIIA, $a_1 \preceq^F_{G_1} a_2 \Leftrightarrow a_2 \preceq^F_{G_1} a_1$, and therefore $a_1 \simeq^F_{G_1} a_2$. By weak transitivity, it is easy to see that $c \prec^F_{G_1} a_1$ and $c \prec^F_{G_1} b$. If we assume $b \preceq^F_{G_1} a_1$, then by weak transitivity, $a_1 \prec^F_{G_1} b$ which contradicts our assumption. So we conclude that $c \prec^F_{G_1} a_1 \prec^F_{G_1} b$.

Now consider the graph G_2 in Figure 5(b). Again, the out-degree requirement of weak transitivity is trivially satisfied, and again by RIIA, $a_1 \simeq^F_{G_2} a_2$. By weak transitivity, it is easy to see that $a_1 \prec^F_{G_2} c$ and $b \prec^F_{G_2} c$. If we assume $a_1 \preceq^F_{G_2} b$, then by weak transitivity, $b \prec^F_{G_2} a_1$ which contradicts our assumption. So we conclude that $b \prec^F_{G_2} a_1 \prec^F_{G_2} c$.

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

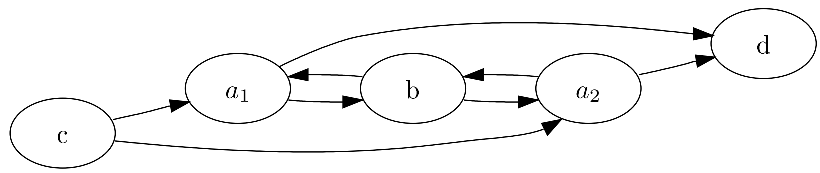

(a) Graph G_1

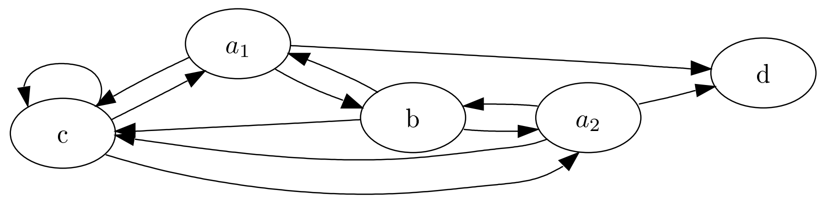

(b) Graph G_2

Figure 5: Graphs for the proof of Theorem 5.1

by RIIA, $a_1 \preceq^F_{G_1} b \Leftrightarrow a_1 \preceq^F_{G_2} b$, which is a contradiction to the fact that $a_1 \prec^F_{G_1} b$ but $b \prec^F_{G_2} a_1$. □

This result is quite a surprise. Intuitively, we would like a ranking procedure to be sensitive to the relative ranking of each agent's voters (transitivity) and not to be influenced by any other seemingly irrelevant information (RIIA). Although these requirements may seem complementary, this impossibility theorem shows that these requirements are in fact contradictory.

If we consider transitivity as a basic requirement, we learn that any axiomatization of a transitive ranking system cannot be restricted to local ordinal properties. That is, when designing a ranking system where transitivity is required, one must choose whether to base the system on some numeric computation, or on ordinal axioms that operate on a global scale.

For example, the standard formalism for the PageRank ranking system in Definition 2.7 and an axiomatization of a similar system suggested by Palacios-Huerta and Volij (2004) are based on numerical computation, while our suggested axiomatization (Altman & Tennenholtz, 2005b) uses ordinal axioms on a global scale. These axioms refer to invariants in relations between ranking of different graphs, rather than between pairs of vertices in the same graph.

The PageRank example demonstrates that some ranking systems may be defined using either of these approaches. We feel that the numeric approach is more suitable for defining and executing ranking systems, while the global ordinal approach is more suitable for axiomatic classification.

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 by Tennenholtz (2004). That algorithm simply starts with ranking by in-degree and refines the ranking as required by strong transitivity until it is satisfied. It is easy to see that the decisions for comparison profiles possible in a 3-vertex graph are dictated by either in-degree or transitivity. Specifically, the profile $\langle(1,3),(2,2)\rangle$ used in the proof above is impossible in such graphs.

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 6a. By weak transitivity, $x_1 \simeq^F_{G_1} x_2 \prec^F_{G_1} b \prec^F_{G_1} a$. (a,b) satisfies the comparison profile $\langle(1,1,2),(3)\rangle$, so we must have $(3) \prec (1,1,2)$. Now let $G_2 \in \mathbb{G}_1$ be the graph in Figure 6b. By weak transitivity $x_1 \simeq^F_{G_2} x_2 \prec^F_{G_2} y \prec^F_{G_2} a \prec^F_{G_2} b$. (b,a) satisfies the comparison profile $\langle(2,3),(1,4)\rangle$, so we must have $(1,4) \prec (2,3)$.

Let $G_3 \in \mathbb{G}_1$ be the graph in Figure 6c. By weak transitivity it is easy to see that $x_1 \simeq^F_{G_3} \cdots \simeq^F_{G_3} x_7 \prec^F_{G_3} y_1 \simeq^F_{G_3} y_2 \prec^F_{G_3} c \prec^F_{G_3} d$. Furthermore, by weak transitivity we conclude that $a \prec^F_{G_3} b$ and $a' \prec^F_{G_3} b'$ from $c \prec^F_{G_3} d$; and $y_1 \prec^F_{G_3} b$ from $x_3 \prec^F_{G_3} d$. Now consider the vertex pair (c,b'). We have shown that $x_1 \simeq^F_{G_3} x_2 \prec^F_{G_3} y_1 \prec^F_{G_3} b$. So, (c,b') satisfies the comparison profile $\langle(1,1,2),(3)\rangle$, thus by RIIA $b' \prec^F_{G_3} c$. Now consider the vertex pair (b,a). We have already shown that $a' \prec^F_{G_3} b' \prec^F_{G_3} c \prec^F_{G_3} d$. So, (a,b) satisfies the comparison profile $\langle(2,3),(1,4)\rangle$, thus by RIIA $b \prec^F_{G_3} a$. However, we have already shown that $a \prec^F_{G_3} b$ – a contradiction. Thus, the ranking system F cannot exist. □

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

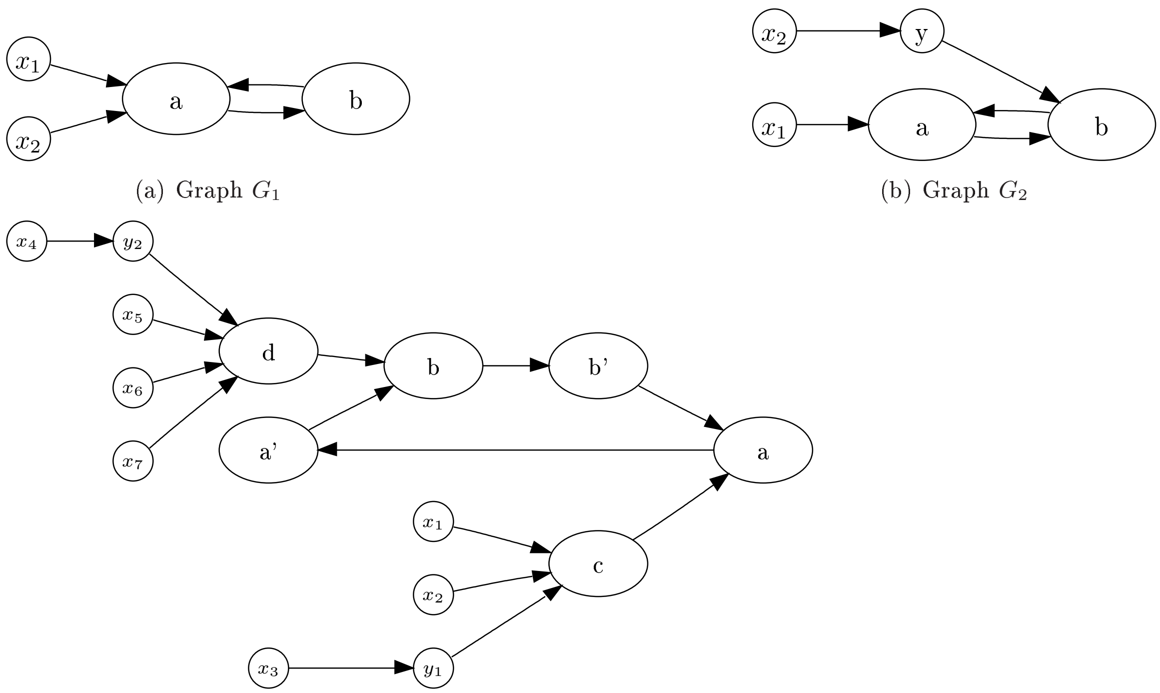

(a) Graph G_1 (b) Graph G_2 (c) Graph G_3

Figure 6: Graphs from the proof of proposition 6.1

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.*

Proof. The proof is exactly the same as for $\mathbb{G}_1$, considering that all graphs in Figure 6 are bipartite. □

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 $\mathbb{G}_{SC}$ that satisfies weak transitivity and RIIA.*

Proof. The proof is similar to the proof of Theorem 5.1, but with an additional vertex e in both graphs that has edges to and from all other vertices. $\square$

7. Relaxing Transitivity

Our impossibility result becomes a possibility result when we relax the transitivity requirement. Instead of comparing only vertices with similar out-degree as in the weak transitivity axiom above, we weaken the requirement for strict preference to hold only in the case where the matching predecessors of one agent are preferred to the *all* predecessors of the other.

Definition 7.1. Let F be a ranking system. We say that F satisfies *strong quasi-transitivity* if for all graphs $G = (V, E)$ and for all vertices $v_1, v_2 \in V$: Assume there is a 1-1 (but not necessarily onto) mapping $f : P(v_1) \mapsto P(v_2)$ s.t. for all $v \in P(v_1)$: $v \preceq f(v)$. Then, $v_1 \preceq v_2$. And, if $P(v_1) \neq \emptyset$ and for all $v \in P(v_1)$: $v \prec f(v)$, then $v_1 \prec v_2$.

Strong quasi transitivity a significantly weaker property than strong transitivity, as it allows for much more indifference in the resulting ranking. Specifically, the ranking system $F_=$ that always ranks all vertices equally satisfies strong quasi transitivity. More generally, any ranking system where the value of a vertex is proportional to a sum of the values of some subset of its predecessors satisfies strong quasi transitivity. We shall see more examples of quasi-transitive ranking systems below.

When we only require strong *quasi*-transitivity and RIIA, we find an interesting family of ranking systems that rank the agents according to their in-degree, breaking ties by comparing the ranks of the strongest predecessors. These recursive in-degree systems work by assigning a rational value for every vertex, that is based on the following idea: rank first based on the in-degree. If there is a tie, rank based on the strongest predecessor's value, and so on. Loops are ranked as periodical rational numbers in base $(n + 1)$ with a period the length of the loop, in the case that continuing on the loop is the maximally ranked option.

The recursive in-degree systems differ in the way different in-degrees are compared. Any monotone increasing mapping of the in-degrees could be used for the initial ranking. To show these systems are well-defined and that the values can be calculated we define these systems algorithmically as follows:

Definition 7.2. Let $r : \mathbb{N} \mapsto \mathbb{N}$ be a monotone nondecreasing function such that $r(i) \leq i$ for all $i \in \mathbb{N}$. The *recursive in-degree ranking system with rank function* r is defined as follows: Given a graph $G = (V, E)$, let $n = |V|$. The relative ranking of two vertices is based on a numeric calculation:

$$v_1 \preceq_G^{RID_r} v_2 \Leftrightarrow \text{value}_r(v_1) \leq \text{value}_r(v_2),$$

where $\text{value}_r(v)$ is defined by maximizing a valuation function $\text{vp}_r(\cdot)$ on all paths that lead to v:

$$\text{value}_r(v) \quad = \quad \max_{\mathbf{a} \in \text{Path}(v)} \text{vp}_r(\mathbf{a}) \tag{1}$$

To ensure the definition is sound, we eliminate loops, and define the path in reverse order:

$$\begin{aligned}\text{Path}(v) \quad = \{ \quad & (v = a_1, a_2, \ldots, a_m) | m \in \mathbb{N}, \\ & (a_m, \ldots, a_1) \text{ is a path in } G \ \wedge (a_{m-1}, \ldots, a_1) \text{ is simple}\}\end{aligned}$$

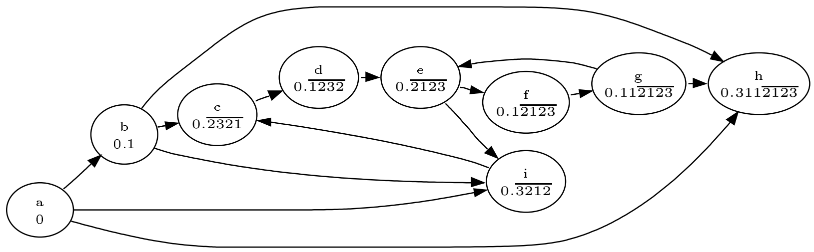

Figure 7: Values assigned by the recursive in-degree algorithm

The path valuation function $\mathrm{vp} : V^* \mapsto \mathbb{Q}$ defines the value to conform to a lexicographic order on in-degrees along the path:

$$\mathrm{vp}_r(a_1, a_2, \ldots, a_m) = \frac{1}{n+1} \left[\begin{array}{l} r(|P(a_1)|) + \\ \begin{cases} 0 & m = 1 \\ \mathrm{vp}_r(a_2, \ldots, a_m, a_2) & a_1 = a_m \wedge m > 1 \\ \mathrm{vp}_r(a_2, \ldots, a_m) & \text{Otherwise.} \end{cases} \end{array} \right] \tag{2}$$

Note that $\mathrm{vp}_r(a_1, a_2, \ldots, a_m)$ is infinitely recursive in the case when the path contains a loop (c.f. $a_1 = a_m \wedge m > 1$). For computation sake we can redefine this case finitely as:

$$\begin{aligned} \mathrm{vp}_r(a_1, \ldots, a_m, a_1) &= \sum_{i=0}^{\infty} \frac{1}{(n+1)^{mi}} \sum_{j=1}^{m} \frac{r(|P(a_j)|)}{(n+1)^j} = \\ &= \frac{(n+1)^m}{(n+1)^m - 1} \mathrm{vp}_r(a_1, \ldots, a_m). \end{aligned} \tag{3}$$

Example 7.3. An example of the values assigned for a particular graph when r is the identity function is given in Figure 7. As $n = 9$, and the definition in (2) is based on recursive division by $n+1$, these values are simply decimals which consist of a concatenation of in-degrees along the maximal path.

The value of zero is assigned to a via the first case in (2), as the only path leading to a consists of a itself. The value for b is arises from the path (b, a) and the third case in (2), the recursive call gives the value of the path (a) which we have seen to be equal 0. This is added to $r(|P(b)|) = 1$ and divided by 10, giving the result 0.1. The values of c, d, e, and i arise from a loop consisting of these vertices. Applying the second case in (2), we have the equations

$$\begin{aligned} \mathrm{value}_r(i) &= \mathrm{vp}_r(i, e, d, c, i) = \frac{1}{10}\left[3 + \mathrm{vp}_r(e, d, c, i, e)\right] \\ \mathrm{value}_r(e) &= \mathrm{vp}_r(e, d, c, i, e) = \frac{1}{10}\left[2 + \mathrm{vp}_r(d, c, i, e, d)\right] \\ \mathrm{value}_r(d) &= \mathrm{vp}_r(d, c, i, e, d) = \frac{1}{10}\left[1 + \mathrm{vp}_r(c, i, e, d, c)\right] \\ \mathrm{value}_r(c) &= \mathrm{vp}_r(c, i, e, d, c) = \frac{1}{10}\left[2 + \mathrm{vp}_r(i, e, d, c, i)\right] \end{aligned}$$

By using (3), we get the periodic decimals seen in Figure 7. The values for vertices f, g, and h are again assigned using the third case in (2). Note that the complete maximal paths to these vertices contain the loop (e, d, c, i, e) and thus all of these vertices' values include a periodic decimal part, as can be seen in Figure 7.

The recursive in-degree system satisfies an interesting fixed point property that can be used to facilitate its efficient computation:

Proposition 7.4. *Let $r : \mathbb{N} \mapsto \mathbb{N}$ be a monotone nondecreasing function such that $r(i) \leq i$ for all $i \in \mathbb{N}$ and define $r(0) = 0$. The value function for the recursive in-degree ranking system satisfies:*

$$value_r(v) \quad = \quad \begin{cases} \frac{1}{n+1}\left[r(|P(v)|) + \max_{p\in P(v)} value_r(p)\right] & P(v) \neq \emptyset \\ 0 & Otherwise \end{cases} \tag{4}$$

Proof. Denote $\mathrm{Path}'(p, v)$ as the set of almost-simple directed paths to p which do not pass through v unless immediately looping back to p:

$$\begin{aligned} \mathrm{Path}'(p, v) \quad = \{ \quad & (p = a_1, a_2, \ldots, a_m)| \\ & (a_m, \ldots, a_1) \text{ is a path in } G \ \wedge (a_{m-1}, \ldots, a_1) \text{ is simple} \wedge \\ & \forall i \in \{1, \ldots, m-2, m\} : a_i \neq v \wedge \\ & a_{m-1} = v \Leftrightarrow a_m = p\}. \end{aligned}$$

Let $v \in V$ be some vertex. Then,

$$\begin{aligned} \mathrm{value}_r(v) \quad &= \quad \max_{\mathbf{a}\in \mathrm{Path}(v)} \mathrm{vp}_r(\mathbf{a}) = \\ &= \quad \frac{1}{n+1}\left[\begin{array}{l} r(|P(v)|) + \max_{(v=a_1,\ldots,a_m)\in \mathrm{Path}(v)} \\ \begin{cases} \mathrm{vp}_r(a_2, \ldots, a_m, a_2) & a_1 = a_m \wedge m > 1 \\ \mathrm{vp}_r(a_2, \ldots, a_m) & \text{Otherwise.} \end{cases} \end{array}\right] = \qquad (5) \\ &= \quad \frac{1}{n+1}\left[r(|P(v)|) + \max_{p\in P(v)} \max_{\mathbf{a}\in \mathrm{Path}'(p,v)} \mathrm{vp}_r(\mathbf{a})\right] = \qquad (6) \\ &= \quad \frac{1}{n+1}\left[r(|P(v)|) + \max_{p\in P(v)} \max_{\mathbf{a}\in \mathrm{Path}(p)} \mathrm{vp}_r(\mathbf{a})\right] = \\ &= \quad \frac{1}{n+1}\left[r(|P(v)|) + \max_{p\in P(v)} \mathrm{value}_r(p)\right]. \end{aligned}$$

Note that (5) is equal to zero 0 if $P(v) = \emptyset$, as required. To show that the equality (6) holds, assume for contradiction that there exists $p \in P(v)$ and $\mathbf{a} \in \mathrm{Path}(p)$ such that

$$\mathrm{vp}_r(\mathbf{a}) > \max_{p'\in P(v)} \max_{\mathbf{a}'\in \mathrm{Path}'(p',v)} \mathrm{vp}_r(\mathbf{a}'). \tag{7}$$

From $\mathbf{a} \in \mathrm{Path}(p) \setminus \mathrm{Path}'(p, v)$, we know that $a_i = v$ for some $i \in \{1, \ldots, m\}$. Assume wlog that i is minimal. Let $\mathbf{b}$ denote the path $(p = a_1, a_2, \ldots, a_i, p)$ and let $\mathbf{c}$ denote the path $(p' = a_{i+1}, \ldots, a_m, a_{j+1}, \ldots, a_{i+1})$ if $a_m = a_j$ for some $j < i$ or $(p' = a_{i+1}, \ldots, a_m)$

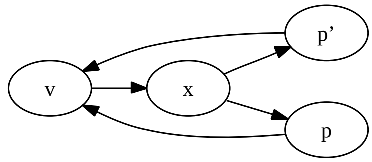

$$\begin{aligned}\mathbf{a} &= (p, x, v, p', x) \\ \mathbf{b} &= (p, x, v, p) \\ \mathbf{c} &= (p', x, v, p')\end{aligned}$$

Figure 8: Example of paths from the proof of Proposition 7.5.

otherwise. An example of such paths is given in Figure 8. Note that $\mathbf{b} \in \mathrm{Path}'(p, v)$ and $\mathbf{c} \in \mathrm{Path}'(p', v)$, where $p, p' \in P(v)$. Now, note that

$$\mathrm{vp}_r(\mathbf{a}) \quad = \quad \frac{(n+1)^j - 1}{(n+1)^j}\mathrm{vp}_r(\mathbf{b}) + \frac{1}{(n+1)^j}\mathrm{vp}_r(\mathbf{c}),$$

and thus $\mathrm{vp}_r(\mathbf{a})$ must be between $\mathrm{vp}_r(\mathbf{b})$ and $\mathrm{vp}_r(\mathbf{c})$, in contradiction to assumption (7). □

Note that although it might look compelling to use this fixed point property as a definition of recursive-indegree, it is not well defined, as loops induce in an infinite series of maximizations the we must prove converges. This is the essence of the proof above. This fixed point property is the basis for the efficient algorithm for recursive-indegree provided below.

We shall now show this ranking system does in fact satisfy RIIA and our weakened version of transitivity.

Proposition 7.5. *Let $r : \mathbb{N} \mapsto \mathbb{N}$ be a monotone nondecreasing function such that $r(i) \leq i$ for all $i \in \mathbb{N}$ and define $r(0) = 0$. The recursive in-degree ranking system with rank function r satisfies strong quasi-transitivity and RIIA.*

Proof. The fixed point result in Proposition 7.4 further implies $0 \leq \mathrm{value}_r(v) < 1$, and thus vertices are ordered first by $r(|P(v)|)$ and then by $\max_{p \in P(v)} \mathrm{value}_r(p)$. Therefore, every comparison profile $\langle \mathbf{a}, \mathbf{b} \rangle$ where $\mathbf{a} = (a_1, \ldots, a_k)$, $\mathbf{b} = (b_1, \ldots, b_l)$ is ranked as follows:

$$f\langle \mathbf{a}, \mathbf{b} \rangle = 1 \quad \Leftrightarrow \quad (k = 0) \vee (r(k) < r(l)) \vee [(r(k) = r(l)) \wedge (a_k \leq b_l)]\,.$$

This ranking of profiles trivially yields strong quasi-transitivity as required. □

We have previously presented a preliminary version of the personalized variant of recursive in-degree (Altman & Tennenholtz, 2006a). The algorithm presented there is based on an equivalent recursive definition for value:

$$\mathrm{value}_r(v) \quad = \quad \mathrm{vp}_r(\mathrm{pv}_r((), v)) \tag{8}$$

$$\mathrm{pv}_r(\mathbf{a}, v) \quad = \quad \begin{cases} (v) & P(v) = \emptyset \\ \left(v, \max_{p \in P(v)} \mathrm{pv}_r(\mathbf{a}, v, p)\right) & v \notin \mathbf{a} \\ (a_k, \ldots, a_m, v) & \mathbf{a} = (a_1, \ldots, a_k = v, \ldots, a_m), \end{cases} \tag{9}$$

where the maximum on the paths is taken over $\mathrm{vp}_r(\mathrm{pv}_r(\mathbf{a}, v, p))$.

Algorithm 1 Efficient algorithm for recursive in-degree

1. Initialize $\text{value}_r(v) \leftarrow \frac{1}{n+1} r(|P(v)|)$ for all $v \in V$, where $r(0)$ is assumed to be 0.
2. Let V' be the set of vertices with incoming edges.
3. Iterate $|V|$ times:
   (a) For every vertex $v \in V'$:
      i. Update $\text{value}_r(v) \leftarrow \frac{1}{n+1}\left[r(|P(v)|) + \max_{p \in P(v)} \text{value}_r(p)\right]$.
4. Sort V' by $\text{value}_r(\cdot)$.
5. Output all vertices in $V \setminus V'$ as weakest, followed by the vertices in V' sorted by $\text{value}_r(\cdot)$ in ascending order.

The fixed point property in (4) satisfies the classical Bellman principle of optimality (Stokey & Lucas, 1989), that is

$$v(x_t) = \max\left[F(x_t, x_{t+1}) + \beta v(x_{t+1})\right].$$

Thus, we can apply a dynamic programming algorithm to efficiently compute these values, as seen in Algorithm 1. Note that due to the limits of the size of the graph we can limit the number of iterations and still ensure an exact result in $O(|V| \cdot |E|)$ time. A simple heuristic for improving the efficiency of the algorithm for practical purposes is to reduce the number of iterations, like in other fixed point algorithms such as PageRank (Page et al., 1998). We shall now prove the correctness and complexity of this algorithm.

Proposition 7.6. *Algorithm 1 outputs vertices in V in the order of $\preceq^{RID}$ as defined in Definition 7.2 and works in $O(|V| \cdot |E|)$ time.*

Proof. Let us first denote

$$\begin{aligned} \text{vp}'_r(a_1, a_2, \ldots, a_m, \ldots) &= \frac{1}{n+1}[r(|P(a_1)| + \text{vp}'_r(a_2, \ldots, a_m, \ldots)] \\ \text{vp}'_r() &= 0. \end{aligned}$$

Note that for all $v \in V$ and for all $a_1, \ldots, a_m \in \text{Path}(v)$: If $a_1, \ldots, a_m$ is simple, $\text{vp}'_r(a_1, \ldots, a_m) = \text{vp}_r(a_1, \ldots, a_m)$. Otherwise if $a_n = a_i$, then $\text{vp}_r(a_1 \ldots, a_m) = \text{vp}'_r(a_1, \ldots a_m, a_{i+1}, \ldots, a_m, \ldots)$. Let $\mathbb{P}(v)$ be the set of all reverse paths to v in G, simple or otherwise. We then have for all $v \in V$:

$$\text{value}_r(v) = \max_{p \in \text{Path}(v)} \text{vp}_r(p) = \max_{p \in \mathbb{P}(v)} \text{vp}'_r(p),$$

because the first loop in $p \in \mathbb{P}(v)$ can be replaced with the one maximizing $\text{vp}_r(\cdot)$, thus increasing value.

The iteration in step 3 of the algorithm calculates for all v:

$$\frac{1}{n+1}\left[r_0 + \max_{p_1\in P(v)}\left[\cdots\frac{1}{n+1}\left[r_{|V|-1} + \max_{p_{|V|}\in P(p_{|V|-1})}\frac{1}{n+1}r_{|V|}\right]\cdots\right]\right],$$

where $r_i = r(|P(p_i)|)$ and $p_0 = v$. This value is equal to

$$\begin{aligned} & \max_{p_1\in P(v)}\max_{p_2\in P(p_1)}\cdots\max_{p_{|V|}\in P(p_{|V|}-1)}\sum_{i=0}^{|V|}\frac{r_i}{(n+1)^{i+1}} = \\ = & \max_{(p_1,\ldots,p_{|V|+1})\in\mathbb{P}_{|V|}(v)}\sum_{i=1}^{|V|+1}\frac{r_i}{(n+1)^i} = \\ = & \max_{p\in\mathbb{P}_{|V|+1}(v)}\mathrm{vp}'_r(v), \end{aligned} \tag{10}$$

where $\mathbb{P}_m(v)$ is the set of all reverse paths of length $\leq m$ to v, simple or otherwise.As there are only $|V|$ vertices, any two vertices that differ in the value assigned by the value function from (1) must also differ the value (10) calculated by the algorithm and in the same direction.

We shall now prove the time complexity of the algorithm, by tracing each step. Steps 1 and 2 take $O(|V|)$ time. The iteration in step 3 is repeated $|V|$ times, and for every vertex in V' performs $O(|P(v)|)$ calculations, so each iteration takes $O(|E|)$ time and thus the total time is $O(|V|\cdot|E|)$. Step 4 takes $O(|V'|\log|V'|) \leq O(|V|\log|E|) \leq O(|V|\cdot|E|)$. Finally, the output step 5 takes $O(|V|)$ time. As every step takes no more than $O(|V|\cdot|E|)$ time, so does the entire algorithm. □

8. Axiomatization of Approval Voting

In Sections 5 and 6 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.

Tennenholtz (2004) has previously shown a non-trivial ranking system that satisfies (weak) transitivity, and in the previous section we have seen such a system for RIIA. However, we have not provided a representation theorem for our new system.

In this section we provide a representation theorem for a ranking system that satisfies RIIA but not weak transitivity — the approval voting ranking system (see Definition 2.5). 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.

Fishburn (1978) has axiomatized the Approval Voting ranking system in the context of social choice, where the output of the algorithm is not a ranking, but rather a set of winners. These two distinct settings are very similar, and thus Fishburn's axiomatization of approval voting is of great relevance to our work. We shall compare these two axiomatizations later in this section.

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

The positive response axiom (sometimes referred to as *positive responsiveness*) essentially means that if an agent receives additional votes, its rank must improve:

Definition 8.1. Let F be a ranking system. F satisfies *positive response* if for all graphs $G = (V, E)$ and for all $(v_1, v_2) \in (V \times V) \setminus E$, $v_1 \neq v_2$, 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 8.2. 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 8.3. 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, $v_1 \preceq_{(V,E)}^F v_2 \Leftrightarrow \pi(v_1) \preceq_{(V,E')}^F \pi(v_2)$.

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

Definition 8.4. 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. 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 8.5. *(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 8.6. *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*

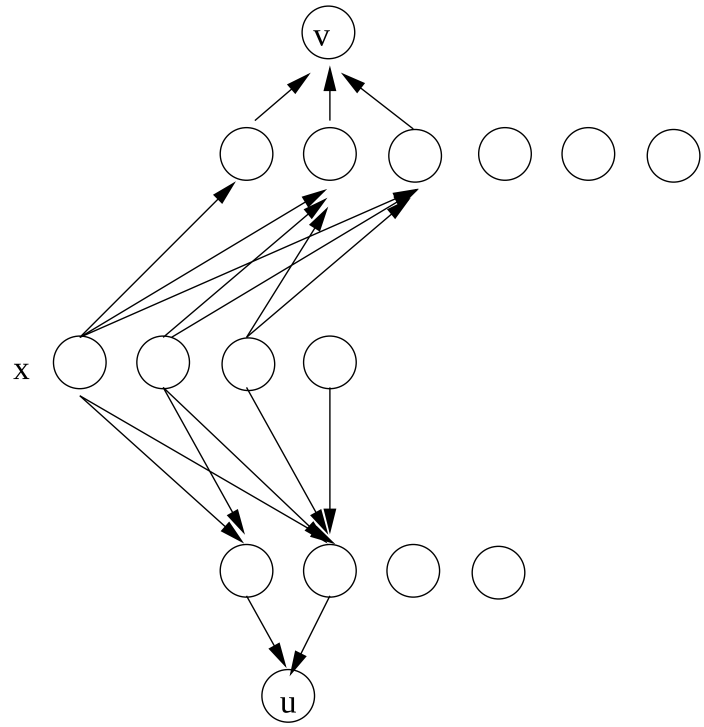

Figure 9: Example of graph G for the profile $\langle(1,3,3),(2,4)\rangle$

Proof. 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 8.5 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.

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 $(\underbrace{1,1,\ldots,1}_{n}) \preceq (\underbrace{1,1,\ldots,1}_{m})$ iff $n \le m$. Let $P = \langle(a_1, \ldots, a_n), (b_1, \ldots, b_m)\rangle$ be a comparison profile. Let $G = (V, E)$ be the following graph (an example of such graph for the profile $\langle(1,3,3),(2,4)\rangle$ is in Figure 9):

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

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 $\langle (\underbrace{1, 1, \ldots, 1}_{n}), (\underbrace{1, 1, \ldots, 1}_{m}) \rangle$, and thus $v \preceq^F_{(V,E')} u \Leftrightarrow n \leq m$.

  By anonymity, $u \preceq^F_{(V,E)} v \Leftrightarrow u \preceq^F_{(V,E')} v$, 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 $\langle (\underbrace{1, 1, \ldots, 1}_{n}), (\underbrace{1, 1, \ldots, 1}_{m}) \rangle$, and thus $v \preceq^F_{(V,E')} u \Leftrightarrow n \leq m$.

  By neutrality, $u \preceq^F_{(V,E)} v \Leftrightarrow u \preceq^F_{(V,E')} v$, 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 $v \preceq^F_{(V,E')} u \Leftrightarrow n \leq m$. Note that $P_G(v) = P_{(V,E')}(v)$ and $P_G(u) = P_{(V,E')}(u)$, so by AIIA, $u \preceq^F_{(V,E)} v \Leftrightarrow u \preceq^F_{(V,E')} v$, and thus as before, $F = AV$.

$\square$

Our axiomatization of approval voting, and specifically the one in (2) above is related to the previous axiomatization by Fishburn (1978). Both axiomatizations share the requirements of Anonymity[1] and Neutrality, but differ in the additional assumptions: Fishburn's requirements refer to relations between the results on different voter sets, which cannot be easily used in the ranking systems setting, as these voters are also alternatives, while our requirements relate to changes in the preferences of a single agent and their ability (positive response) or inability (AIIA) to influence the final result. Our requirements may be mapped to Fishburn's setting and would probably lead to a distinct axiomatization of approval voting in that setting.

9. Concluding Remarks

Reasoning about preferences and preference aggregation is a fundamental task in reasoning about multi-agent systems (see e.g. Boutilier, Brafman, Domshlak, Hoos, & Poole, 2004; Conitzer & Sandholm, 2002; LaMura & 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 PageRank and eBay's reputation systems, among many others).

1. Fishburn does not consider Anonymity as an axiom, but rather defines his social choice model to allow only for anonymous functions.

Moreover, the task of building successful and effective on-line trading environments has become a central challenge to the AI community (Boutilier, Shoham, & Wellman, 1997; Monderer, Tennenholtz, & Varian, 2000; Sandholm, 2003). Ranking systems are believed to be fundamental for the establishment of such environments. Although reputation has always been a major issue in economics (see e.g. Kreps & Wilson, 1982; Milgrom & Roberts, 1982), reputation systems have become so central recently due to the fact that some of the most influential and powerful Internet sites and companies have put reputation systems in the core of their business.

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 (Gibbard, 1973; Satterthwaite, 1975) 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 & Wellman, 1989) and multi-agent coordination (Kfir-Dahav & 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, complemented by a new ranking algorithm, and a representation theorem for the well-known approval voting scheme.

Acknowledgements

This work has been partially supported by a grant from the Israeli Science Foundations (ISF).

References

Altman, A., & Tennenholtz, M. (2005). Ranking systems: the PageRank axioms. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pp. 1–8, New York, NY, USA. ACM Press.

Altman, A., & Tennenholtz, M. (2006). Quantifying incentive compatibility of ranking systems.. In *Proc. of AAAI-06*.

Altman, A., & Tennenholtz, M. (2007a). An axiomatic approach to personalized ranking systems. In *Proc. 20th International Joint Conference on Artificial Intelligence*.

Altman, A., & Tennenholtz, M. (2007b). Incentive compatible ranking systems. In *Proc. of AAMAS-07*.

Arrow, K. (1963). *Social Choice and Individual Values (2nd Ed.)*. Yale University Press.

Bogomolnaia, A., Moulin, H., & Stong, R. (2005). Collective choice under dichotomous preferences. *Journal of Economic Theory*, *122*(2), 165–184. available at http://ideas.repec.org/a/eee/jetheo/v122y2005i2p165-184.html.

Boutilier, C., Shoham, Y., & Wellman, M. (1997). Special issue on economic principles of multi-agent systems. *Artificial Intelligence*, *94*.

Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., & Poole, D. (2004). Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements.. *J. Artif. Intell. Res. (JAIR)*, *21*, 135–191.

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

Conitzer, V., & Sandholm, T. (2002). Complexity of mechanism design. In *Proceedings of the 18th conference on uncertainity in Artificial Intelligence (UAI-02)*, pp. 103–110.

Doyle, J., & Wellman, M. (1989). Impediments to Universal Preference-Based Default Theories. In *Proceedings of the 1st conference on principles of knowledge representation and reasoning*.

Fishburn, P. C. (1978). Axioms for approval voting: Direct proof. *Journal of Economic Theory*, *19*(1), 180–185. available at http://ideas.repec.org/a/eee/jetheo/v19y1978i1p180-185.html.

Gibbard, A. (1973). Manipulation of voting schemes. *Econometrica*, *41*, 587–601.

Kfir-Dahav, N. E., & Tennenholtz, M. (1996). Multi-Agent Belief Revision. In *Proceedings of the 6th conference on theoretical aspects of rationality and knowledge (TARK)*.

Kreps, D., & Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory*, *27*, 253–279.

LaMura, P., & Shoham, Y. (1998). Conditional, Hierarchical Multi-Agent Preferences. In *Proceedings of Theoretical Aspects of Rationality and Knowledge*, pp. 215–224.

May, K. O. (1952). A set of independent, necessary and sufficient conditions for simple majority decision. *Econometrica*, *20*(4), 680–84.

Milgrom, P., & Roberts, J. (1982). Predation, reputation and entry deterrence. *Journal of Economic Theory*, *27*, 280–312.

Monderer, D., Tennenholtz, M., & Varian, H. (2000). Game theory and artificial intelligence. Special issue of Games and Economic behavior.

Moulin, H. (1991). *Axioms of Cooperative Decision Making*. Cambridge University Press.

Page, L., Brin, S., Motwani, R., & Winograd, T. (1998). The PageRank citation ranking: Bringing order to the web. Technical Report, Stanford University.

Palacios-Huerta, I., & Volij, O. (2004). The measurement of intellectual influence. *Econometrica*, *73*(3).

Resnick, P., & Zeckhauser, R. (2001). 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.

Resnick, P., Zeckhauser, R., Friedman, R., & Kuwabara, E. (2000). Reputation systems. *Communications of the ACM*, *43*(12), 45–48.

Sandholm, T. (2003). Making markets and democracy work: A story of incentives and computing. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-03)*, pp. 1649–1671.

Satterthwaite, M. (1975). Strategy proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions.. *Journal of Economic Theory*, *10*, 187–217.

Stokey, N. L., & Lucas, R. E. (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

Tennenholtz, M. (2004). Reputation systems: An axiomatic approach. In *Proceedings of the 20th conference on uncertainity in Artificial Intelligence (UAI-04)*.