Quantifying Incentive Compatibility of Ranking Systems

Alon Altman   Moshe Tennenholtz

15th March 2006

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. Such ranking systems do not exist in empty space, and therefore agents' incentives should be carefully considered. In this paper we define three measures for quantifying the incentive compatibility of ranking systems. We apply these measures to several known ranking systems, such as PageRank, and prove tight bounds on the level of incentive compatibility under two basic properties: strong monotonicity and non-imposition. We also introduce two novel non-imposing ranking systems, one general, and the other for the case of systems with three participants. A full axiomatization is provided for the latter.

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 algorithmPage *et al.* (1998) and eBay's reputation systemResnick & Zeckhauser (2001).

The ranking systems setting can be viewed as a variation of the classical theory of social choiceArrow (1963), where the set of agents and the set of alternative *coincide*. Specifically, we consider dichotomous ranking systems, in which the agents vote for a subset of the rest of the agents. This is a natural representation of the web page ranking settingTennenholtz (2004), where the Internet pages are represented by the agents/alternatives, and the links are represented by votes.

Some basic work targeted at the foundations of ranking systems has been recently initiated. In particular, basic properties of ranking systems have been shown to be impossible to simultaneously accommodateAltman & Tennenholtz (2005a), various known ranking systems have been recently compared with regard to certain criteria by Borodin *et al.* (2005), and several ranking rules have been axiomatized Altman & Tennenholtz (2005b); Palacios-Huerta & Volij (2004); Slutzki & Volij (2005).

Although the above mentioned work consists of a significant body of rigorous research on ranking systems, the study did not consider the effects of the agents' incentives on ranking systems. The issue of incentives has been extensively studied in the classical social choice literature. The Gibbard–Satterthwaite theorem (see Mas-Colell, Whinston, & Green (1995)) shows that in the classical social welfare setting, it is impossible to aggregate the rankings in a strategy-proof fashion under some basic conditions. The incentives of the candidates themselves were considered in the context of electionsDutta, Jackson, & Le Breton (2001), where a related impossibility result is presented. Another notion of incentives was considered in the case where a single agent may create duplicates of itselfCheng & Friedman (2005).

In the context of ranking systems, Altman & Tennenholtz (2006) have defined a notion of a (fully) incentive compatible ranking system and have categorized the incentive compatible ranking systems satisfying several basic properties. In this paper, we generalize this notion of incentive compatibility to systems that allow deviations to a limited degree and provide tight bounds for the levels of incentive compatibility under two important properties, suggesting and evaluating practical ranking systems in the process.

We define three notions of limited incentive compatibility and use these notions to quantify the incentive compatibility of ranking systems. Specifically, we quantify the incentive compatibility of the Approval Voting and PageRank ranking systems and prove a significant lower bound on the incentive compatibility of any ranking system satisfying the basic *strong monotonicity* property, which is satisfied by almost all practical ranking systems.

In addition, we consider *non imposing* ranking systems, i.e. systems in which any strict ordering of the agents is feasible. We show that although there are no fully incentive compatible non imposing ranking systems, we can find such a non imposing ranking system that is incentive compatible up to a deviation by one agent by at most one rank.

Finally, we provide a full axiomatization of a non-imposing incentive compatible ranking system for the setting with exactly three agents.

2 Ranking systems and Incentive Compatibility

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 on a vertex set V that do not include self edges. A *ranking system* F is a functional that for every finite vertex set V maps graphs $G \in \mathbb{G}_V$ to an ordering $\preceq^F_G \in L(V)$. If F is defined only on a subset of $\mathbb{G}_V$ we call it a *partial ranking system.*

One can view this setting as a variation/extension of the classical theory of social choice as modeled by Arrow (1963). The vertices in the ranking systems setting correspond to the agents and alternatives in the social choice setting, and the edges correspond to the votes. In the sequel, we will use these terms interchangeably. The ranking systems setting differs from the classical social choice setting 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.

2.1 Basic Properties of Ranking Systems

Now we define some basic properties of ranking systems to guide our quantification.

A basic requirement from a ranking system is that when there are no votes (or all votes) in the system, all agents must be ranked equally. We call this requirement minimal fairness.

Definition 2.4. A ranking system F is *minimally fair* if for every graph $G_\perp = (V, \emptyset)$ with no edges, and for every graph $G_\top = (V, V \times V \setminus \{(v,v)|v \in V\})$ with all edges and for every $v_1, v_2 \in V$: $v_1 \simeq^F_{G_\perp} v_2$ and $v_1 \simeq^F_{G_\top} v_2$.

Another basic requirement from a ranking system is that as agents gain additional votes, their rank must improve, or at least not worsen. Surprisingly, this vague notion can be formalized in (at least) two distinct ways: The monotonicity property considers the situation where one agent has a superset of the votes another agent has *in the same graph*, where the positive response property considers the addition of a vote for an agent *between graphs*. This distinction is important because the two properties are neither equivalent, nor imply each other.

Notation 2.5. Let $G = (V, E)$ be a graph, and let $v \in V$ be a vertex. The predecessor set of v is $P_G(v) = \{v'|(v', v) \in E\}$. The successor set of v is $S_G(v) = \{v'|(v, v') \in E\}$.

Definition 2.6. Let F be a ranking system. F satisfies *weak 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 $G' = (V, E \cup (v_1, v_2))$. Then, $v_3 \preceq^F_G v_2$ implies $v_3 \preceq^F_{G'} v_2$ and $v_3 \prec^F_G v_2$ implies $v_3 \prec^F_{G'} v_2$.

Definition 2.7. A ranking system F satisfies *weak monotonicity* if for all $G = (V, E)$ and for all $v_1, v_2 \in V$: If $P(v_1) \subseteq P(v_2)$ then $v_1 \preceq^F_G v_2$. F furthermore satisfies *strong monotonicity* if $P(v_1) \subsetneq P(v_2)$ additionally implies $v_1 \prec^F_G v_2$.

Another basic property of ranking system is that all strict rankings must have a graph that generates them.

Definition 2.8. Let F be a ranking system, F satisfies *non imposition* if for all V and for all strict linear orderings $L \in L(V)$: there exists some $G \in \mathbb{G}_V$ such that $F(G) \equiv L$.

The aforementioned conditions are basic in the sense that all well-known ranking systems such as PageRank, Hubs&Authorities, and Approval Voting satisfy all of them.

2.2 Incentive Compatibility

Ranking systems do not exist in empty space. The results given by ranking systems frequently have implications for the agents being ranked, which are the same agents that determine in the ranking. Therefore, the incentives of these agents should in many cases be taken into consideration.

In our approach, we aim that a ranking system will minimize agents' gain of rank for stating untrue preferences, under the assumption that the agents are interested only in their own ranking (and not, say, in the ranking of those they prefer). We further assume that an agent in interested in its *expected* rank, assuming equally ranked agents are ordered randomly. Formally, the expected rank (henceforth referred to simply as *rank*) is defined as follows:

Definition 2.9. The *rank* of a vertex v in a graph G under the ranking system F is defined as

$$\begin{aligned} r_G^F(v) &= |\{v' : v' \prec v\}|\,| + \tfrac{1}{2}\,|\{v' : v' \simeq v\}| = \\ &= \tfrac{1}{2}\,|\{v' : v' \prec v\}| + \tfrac{1}{2}\,|\{v' : v' \preceq v\}|\,. \end{aligned}$$

Given this definition of rank, we can now define the magnitude of an agent's best deviation:

Definition 2.10. Let F be ranking system and let $G = (V, E)$ be a graph for which F is defined. The *deviation magnitude* $\delta_G^F(v)$ *of* v *in* G *under ranking system* F is defined as $\max\{\, r^F_{(V,E')}(v) - r^F_G(v) \Big|\, F(V, E')$ is defined, $\forall v' \in V \setminus \{v\}, v'' \in V : (v', v'') \in E \Leftrightarrow (v', v'') \in E'\}$. That is, the maximum rank difference v can obtain for itself by changing its outgoing vertices in G under F.

Now we can consider the case where no agent ever has a useful deviation:

Definition 2.11. Let F be ranking system. F is called *(fully) incentive compatible* over a set of graphs $\mathbb{G}$ if for all graphs $G \in \mathbb{G}$ and for all $v \in V$: $\delta_G^F(v) = 0$.

It has been shown Altman & Tennenholtz (2006) that there are no fully incentive compatible ranking systems satisfying all of the basic properties we have outlined above. Therefore, it is essential to weaken this requirement of incentive compatibility. This can be done in three different ways:

Definition 2.12. Let F be a ranking system. F is called k*-worst case incentive compatible* over a set of graphs $\mathbb{G}$ if for all graphs $G \in \mathbb{G}$ and for all $v \in V$: $\delta_G^F(v) \leq k$. We say that the *worst case incentive compatibility* of F is k if it is k-worst case incentive compatible, but not $(k - \varepsilon)$-worst case incentive compatible for all $\varepsilon > 0$.

Definition 2.13. Let F be a ranking system. F is called k-*mean incentive compatible* over a set of graphs $\mathbb{G}$ if for all graphs $G \in \mathbb{G}$: $\sum_{v \in V} \delta_G^F(v)/|V| \leq k$. We say that the *mean incentive compatibility* of F is k if it is k-mean incentive compatible, but not $(k-\varepsilon)$-mean incentive compatible for all $\varepsilon > 0$.

Definition 2.14. Let F be a ranking system. F is called k-*agent incentive compatible* for a set of graphs $\mathbb{G}$ if for all graphs $G \in \mathbb{G}$: $|\{v \in V | \delta_G^F(v) > 0\}| \leq k$. We say that the *agent incentive compatibility* of F is k if it is k-agent incentive compatible, but not $(k-1)$-agent incentive compatible.

Note that when k is zero, all of these definitions coincide with full incentive compatibility as defined above.

Of the basic properties we defined above, Altman & Tennenholtz (2006) have shown that weak positive response, weak monotonicity and minimal fairness could each be satisfied by an (artificially constructed) fully incentive compatible ranking system. This leads us to concentrate on the levels of incentive compatibility attainable under strong monotonicity and non-imposition. In these fundamental cases, full incentive compatibility cannot be obtained, and thus it is interesting to try and obtain a more limited degree of incentive compatibility. In the sequel we show tight bounds for the levels of incentive compatibility under these two conditions.

3 Incentive Compatibility Under Strong Monotonicity

When we study the incentive compatibility of ranking systems satisfying strong monotonicity, it is helpful to keep in mind that this property is satisfied by almost all practical ranking systems, including Approval Voting, PageRank, and Hubs&Authorities. Specifically, we are going to quantify the incentive compatibility of the Approval Voting and PageRank ranking systems, when the out-degree of each vertex is limited to some constant k.

Definition 3.1. The *approval voting* ranking system AV is defined as $a \preceq^{AV} b \Leftrightarrow |P(a)| \leq |P(b)|$.

First, we are going to prove a general negative result about ranking systems that satisfy strong monotonicity.

Theorem 3.2. *There exists no strongly monotone ranking system that is $(\frac{k}{2} - \varepsilon)$-worst case incentive compatible on the set of graphs with max out-degree k for all $\varepsilon > 0$. Furthermore, there exists no minimally fair strongly monotone ranking system that is $(\frac{k}{2} - \varepsilon)$-mean incentive compatible on the set of graphs with max out-degree k for all $\varepsilon > 0$.*

Proof. Assume a strongly monotone ranking system F and assume a graph $G = (V, E)$ with $k+1$ vertices $V = \{v_0, v_1, \ldots, v_k\}$ and edges $E = \{(v_0, v_1), (v_0, v_2), \ldots, (v_0, v_k)\}$. Assume a strongly monotone l-worst case incentive compatible ranking system F. By strong monotonicity, F ranks

$$v_0 \prec_G^F v_1 \simeq_G^F v_2 \simeq_G^F \cdots \simeq_G^F v_k.$$

This gives $r^F_G(v_0) = \frac{1}{2}$. However, if v_0 changes its votes to $\emptyset$, the rank will become (by strong monotonicity) $v_0 \simeq v_1 \simeq v_2 \simeq \cdots \simeq v_k$, and thus $r^F_{G'}(v_0) = \frac{k+1}{2}$. We have shown a manipulation of magnitude $\frac{k}{2}$, in contradiction to the fact that F is $(\frac{k}{2} - \varepsilon)$-IC, where $\varepsilon > 0$.

Now assume a minimally fair strongly monotone ranking system F'. We will show a graph $G = (V, E)$ in which all agents have a deviation of magnitude $\frac{k}{2}$. The graph is the complete clique with $k+1$ vertices: $V = \{0, \ldots, k\}$ and $E = V \times V \setminus \{(v, v)|v \in V\}$. Note that this graph has a maximal out-degree of k and F ranks all agents equally (due to minimal fairness). However, if any agent v removes all its outgoing edges to form a graph G', then that agent will be, by strong monotonicity, ranked above all other agents. Thus, $r^{F'}_G(v) = \frac{k+1}{2}$, while $r^{F'}_{G'}(v) = k + \frac{1}{2}$. Thus $\delta^{F'}_G(v) = \frac{k}{2}$ for all $v \in V$. Therefore, F' is not $(\frac{k}{2} - \varepsilon)$-mean incentive compatible for all $\varepsilon > 0$. □

We can now quantify the incentive compatibility of the approval voting ranking system, showing that the aforementioned lower bound is tight.

Proposition 3.3. *The approval voting ranking system AV satisfies the following over the set of graphs with max out-degree k:*

- *The worst case incentive compatibility of AV is $\frac{k}{2}$.*
- *The mean incentive compatibility of AV is $\frac{k}{2}$.*
- *The agent incentive compatibility of AV over the set of graphs with n vertices ($n > 1$) is n.*

Proof. First we will prove that AV_k is $\frac{k}{2}$-worst case incentive compatible. Let $G = (V, E), G' = (V, E') \in \mathbb{G}$ be graphs that differ in the outgoing edges from v. Note that $|P_G(v)| = |P_{G'}(v)|$ as neither G nor G' include self-edges. Let $S_{del} = \{u \in S_G(v) \setminus S_{G'}(v)\}$. Note that $|S_{del}| \le k$. For all $u \in V \setminus S_{del}$: $|P_G(u)| = |\{w|(w, u) \in E\}| \le |\{w|(w, u) \in E'\}| = |P_{G'}(u)|$, and thus $|P_{G'}(u)| < |P_{G'}(v)| \Rightarrow |P_G(u)| < |P_G(v)|$ and $|P_{G'}(u)| \le |P_{G'}(v)| \Rightarrow |P_G(u)| \le |P_G(v)|$. Furthermore, for all $u \in S_{del}$: $|P_{G'}(u)| = |P_G(u)| + 1$. Let $S_a = \{u \in S_{del} : |P_G(u)| = |P_G(v)|\}$ and $S_b = \{u \in S_{del} : |P_G(u)| + 1 = |P_G(v)|\}$ Now,

$$\begin{aligned} r^{AV}_{G'}(v) - r^{AV}_G(v) &= \tfrac{1}{2}\,|\{v'|v' \prec_{G'} v\}| \\ &\quad -\tfrac{1}{2}\,|\{v'|v' \prec_G v\}| \\ &\quad +\tfrac{1}{2}\,|\{v'|v' \preceq_{G'} v\}| \\ &\quad -\tfrac{1}{2}\,|\{v'|v' \preceq_G v\}| \\ &\le \tfrac{1}{2}|S_a| + \tfrac{1}{2}|S_b| \le \tfrac{1}{2}|S_{del}| \le \tfrac{k}{2}. \end{aligned}$$

The $\frac{k}{2}$-mean incentive compatibility immediately follows, and the n-agent incentive compatibility is trivial. AV_k satisfies strong monotonicity and minimal fairness, and thus it is not $(\frac{k}{2} - \varepsilon)$-mean incentive compatible, and not $(\frac{k}{2} - \varepsilon)$-worst case incentive compatible for all $\varepsilon > 0$.

To show that AV_k is not $(n-1)$-agent incentive compatible over the set of graphs with n vertices ($n > 1$), assume the full loop with n vertices $G = (V, E)$ defined as follows:

$$\begin{aligned} V &= \{0, \ldots, n-1\} \\ E &= \{(i, i+1 \mod n), | i = 0 \ldots n-1\} \end{aligned}$$

Now, by removing all of its edges, each agent can improve its own relative rank by $\frac{1}{2}$, and thus all n agents have a deviation, and thus AV is not $(n-1)$-agent incentive compatible. □

We now define the PageRank matrix which is the matrix which captures the random walk created by the PageRank procedure. 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.

Definition 3.4. Let $G = (V, E)$ be a directed 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 3.5. Let $G = (V, E)$ be some strongly connected graph, and assume $V = \{v_1, v_2, \ldots, v_n\}$. Let $0 < d < 1$ be a *damping factor*. Let $\vec{r}$ be the unique solution of the system $(1-d) \cdot A_G \cdot \vec{r} + d \cdot \begin{pmatrix} 1 & 1 & \cdots & 1 \end{pmatrix}^T = \vec{r}$ where $\sum r_i = n$. 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 $PR_G(v_i) \leq PR_G(v_j)$.

We will now quantify the incentive compatibility of the PageRank ranking system:

Proposition 3.6. *The PageRank ranking system PR with damping factor d is not $(\frac{n}{2} - 2)$-mean incentive compatible nor $(n-1)$-agent incentive compatible on the set of graphs with n vertices ($n > 2$) and out-degree 1.*

Proof. Consider the graph $G = (V, E)$ where $V = \{0, \ldots, n-1\}$ and $E = \{(i, i+1 \mod n) | i = 0, \ldots n-1\}$. In this graph PR ranks all agents equally due to symmetry. Let $v \in V$ be some agent. Assume wlog $v = n-1$ and let $G' = (V, E')$ be defined as $E' = E \setminus \{(n-1, 0)\} \cup \{(n-1, n-2)\}$. Applying linear algebra, we conclude that PR ranks $0 \prec 1 \prec \cdots \prec n-3 \prec n-1 \prec n-2$ in G' and thus $r_{G'}^{PR}(v) = r_{G'}^{PR}(n-1) = n - 1.5$. However, $r_G^{PR}(v) = \frac{n}{2}$, and thus $\delta_G^{PR}(v) \geq \frac{n-3}{2}$. This is true for all $v \in V$, so we see that PR is not $(\frac{n}{2} - 2)$-mean incentive compatible nor $(n-1)$-agent incentive compatible for an arbitrary graph G with n vertices. □

A similar lower bound showing deviations of magnitude $O(n)$ by all agents can be shown for the Hubs&Authorities ranking system as presented by Kleinberg (1999).

4 Non-imposing Ranking Systems

Recall that non-imposing ranking systems are those that accommodate any strict order on the vertices. We will now show that this requirement cannot be satisfied when requiring full incentive compatibility.

Fact 4.1. *There exists no non-imposing incentive compatible ranking system.*

Proof. Assume the vertex set $V = \{v_1, v_2\}$. There are two potential edges in this graph $e_1 = (v_1, v_2)$ and $e_2 = (v_2, v_1)$. Let $G = (V, E)$ be a graph s.t. $v_1 \prec_G v_2$ and let G' be a graph s.t. $v_2 \prec_{G'} v_1$. As $r_{G'}(v_1) \neq r_G(v_1)$ and $r_G(v_2) \neq r_{G'}(v_2)$, from incentive compatibility, the symmetric difference $E \oplus E' = (E \cup E') \setminus (E \cap E') = \{e_1, e_2\}$. Let $E'' = E \oplus \{e_1\} = E' \oplus \{e_2\}$. From incentive compatibility $r_{G''}(v_1) = r_G(v_1) = \frac{1}{2} = r_{G'}(v_2) = r_{G''}(v_2)$, but this cannot be as if $v_1 \simeq_{G''} v_2$, $r_{G''}(v_1) = r_{G''}(v_1) = 1$. ☐

We will now show a 1-worst case incentive compatible ranking system satisfying non-imposition. This ranking system is also 1-agent incentive compatible, which sets a tight bound.

Theorem 4.2. *There exists a ranking system F that satisfies non-imposition, 1-worst case incentive compatibility, $\frac{1}{n}$-mean incentive compatibility on graphs with n vertices, 1-agent incentive compatibility, and weak positive response.*

Proof. The ranking system F is defined as follows: Assume a graph $G = (V, E)$ with $V = \{v_1, v_2, \ldots, v_n\}$. For each $v \neq v_1$ we define

$$p(v) = \begin{cases} |P(v) \setminus S(v_1)| + n & v \in S(v_1) \\ |P(v) \cap S(v_1)| & v \notin S(v_1) \end{cases}.$$

Now we define a strict ordering $\preceq^*$ on $V \setminus \{v_1\}$:

$$\begin{aligned} v_i \preceq^* v_j \quad \Leftrightarrow \quad & [p(v_i) < p(v_j)] \vee \\ & \vee [p(v_i) = p(v_j) \wedge i \leq j]. \end{aligned}$$

Given this ordering we can finally define $\preceq_G^F$:

$$\begin{aligned} v_i \preceq_G^F v_j \quad \Leftrightarrow \quad & (i \neq 1 \wedge j \neq 1 \wedge v_i \preceq^* v_j) \vee \\ & \vee (i = 1 \wedge |\{u | u \preceq^* v_j\}| \geq |P(v_1)|) \vee \\ & \vee (j = 1 \wedge |\{u | u \preceq^* v_i\}| < |P(v_1)|). \end{aligned}$$

The weak positive response property is satisfied because addition of an edge (u, v) either weakly increases $p(v)$ if $v \neq v_1$, increasing the relative rank of v, or increases $|P(v_1)|$ if $v = v_1$, and thus again increases the relative rank of v.

To prove F satisfies non-imposition, assume a vertex set $V = \{v_1, \ldots, v_n\}$ and strict ordering $\preceq'$ on V. Let $u_1, u_2, \ldots, u_{n-1}$ be the vertices in $V \setminus \{v_1\}$ ordered according to $\preceq'$ and let $k = |\{v \in V : v \preceq' v_1\}|$. Let $G = (V, E)$ be the graph defined as follows:

$$\begin{aligned} E &= \{(v_1, u_i) | i > \tfrac{n-1}{2}\} \cup \{(u_i, v_1) | i < k\} \cup \\ &\quad \{(u_i, u_j) | \tfrac{n-1}{2} < i < j - 1 + \tfrac{n}{2}\} \cup \\ &\quad \{(u_i, u_j) | i < j - \tfrac{n-1}{2}\}. \end{aligned}$$

First note that for all $u_i \in V \setminus \{v_1\}$:

$$p(u_i) = \begin{cases} i + \lfloor \frac{n}{2} \rfloor & i > \frac{n-1}{2} \\ i - 1 & \text{Otherwise} \end{cases}$$

Thus, $u_1 \prec^* u_2 \prec^* \cdots \prec^* u_{n-1}$. As $|P(v_1)| = k - 1$, $u_1 \prec^F_G \cdots \prec^F_G u_{k-1} \prec^F_G v_1 \prec^F_G u_k \prec^F_G \cdots \prec^F_G u_{n-1}$, and thus $\preceq^F_G \equiv \preceq'$, as required.

We will now prove the incentive compatibility features of this ranking system. Let $G = (V, E)$ be some graph. Note that both $\preceq^*$ and $\preceq^F_G$ are strict orderings. The deviation magnitude of agent v_1 is 0, as its rank is dependent only on its in-degree, which it cannot manipulate:

$$\begin{aligned} \delta^F_G(v_1) &= \max\{r^F_{(V,E')}(v) - r^F_G(v)\} = \\ &= \max\{(|P_{(V,E')}(v_1)| + \tfrac{1}{2}) - (|P_G(v_1)| + \tfrac{1}{2})\} = \\ &= \max\{|P_G(v_1)| - |P_G(v_1)|\} = 0. \end{aligned}$$

Let $v_i \in V \setminus \{v_1\}$ be an agent. The rank $r^F_G(v_i)$ is:

$$\begin{aligned} r^F_G(v_j) &= \tfrac{1}{2}\left|\{v' : v' \prec^F_G v_i\}\right| + \tfrac{1}{2}\left|\{v' : v' \preceq^F_G v_i\}\right| = \\ &= \left|\{v' : v' \prec^F_G v_i\}\right| + \tfrac{1}{2} = \\ &= \begin{cases} |\{v' : v' \prec^* v_i\}| + 1.5 & |\{v' | v' \preceq^* v_i\}| \geq |P(v_1)| \\ |\{v' : v' \prec^* v_i\}| + \frac{1}{2} & \text{Otherwise.} \end{cases} \end{aligned}$$

Now, $|\{v' : v' \prec^* v_i\}|$ is independent of the outgoing edges of v_i given $S(v_1)$, as $v_i \in S(v_1)$ iff its outgoing edges are used to rank agents $\notin S(v_1)$. Thus, the only manipulation v_i might do is to change $|P(v_1)|$, and thus increase its rank by 1. In order to increase its rank, v_i must decrease $|P(v_1)|$. v_i can do so by at most 1, by removing an edge (v_i, v_1) if it exists. This manipulation can only be done if $|\{v' | v' \preceq^* v_i\}| = |P_G(v_1)|$. As $\preceq^*$ is strict and $0 \leq |P_G(v_1)| \leq n - 1$, there exists exactly one agent v_i satisfying this condition.

Thus, for some $v_i \in V$: $\delta^F_G(v_i) \leq 1$, and for all $v_j \in V \setminus \{v_i\}$: $\delta^F_G(v_i) = 0$. So we conclude that F is 1-worst case incentive compatible, $\frac{1}{n}$-mean incentive compatible on graphs with n vertices, and 1-agent incentive compatible. □

4.1 A Fully Incentive Compatible Non-imposing Ranking System for 3 Agents

We have previously shown that there exists no general incentive compatible non-imposing ranking system. However, if we limit our domain we may find that there exist such

		$v_0 \to v_1$	$v_0 \to v_2$
$v_2 \to v_0$	$v_1 \to v_2$	$\simeq$	$v_1 \prec v_0 \prec v_2$
	$v_1 \to v_0$	$v_2 \prec v_1 \prec v_0$	$v_1 \prec v_2 \prec v_0$
$v_2 \to v_1$	$v_1 \to v_2$	$v_0 \prec v_2 \prec v_1$	$v_0 \prec v_1 \prec v_2$
	$v_1 \to v_0$	$v_2 \prec v_0 \prec v_1$	$\simeq$

Figure 1: Schematic representation of the three-plurality ranking systems

ranking systems. In this section, we will provide a full axiomatization for non-imposing incentive compatible ranking systems when there are exactly three agents.

Definition 4.3. A ranking system is called *three-plurality* if for every graph $G = (V, E)$ such that $|V| = 3$ there exists an ordering v_0, v_1, v_2 of the vertices in V such that F ranks $u \preceq v \Leftrightarrow f(u) \le f(v)$, where $f(v)$ is one of the following:

$$\begin{aligned} f_1(v_i) &= I\left[(v_{i-1}, v_i) \in E\right] + I\left[(v_{i+1}, v_{i-1}) \notin E\right] \\ f_2(v_i) &= I\left[(v_{i-1}, v_i) \in E \wedge (v_{i-1}, v_{i+1}) \notin E\right] + \\ &\quad + I\left[(v_{i+1}, v_i) \in E \vee (v_{i+1}, v_{i-1}) \notin E\right], \end{aligned}$$

where all the indices are calculated modulo 3, and I is the indicator function.

There are exactly four three-plurality ranking systems for graphs with $V = \{v_0, v_1, v_2\}$. These ranking systems all implement plurality voting when each agent must vote, as illustrated in Figure 1, and differ in the interpretation of the cases where agents cast no votes or both votes.

Theorem 4.4. *Let F be a ranking system over the set of graphs with 3 vertices. F is three-plurality iff it satisfies all of the following criteria: incentive compatibility, non-imposition, weak positive response, and minimal fairness.*

Furthermore, these conditions are independent.

Proof. We must first show that any three-plurality ranking system F satisfies these four criteria. Incentive compatibility and non-imposition can easily be deduced from Figure 1. To show that F satisfies weak positive response, note that any added edge (v_i, v_j) may only increase $f(v_j)$ and decrease $f(v_k)$ for $k \neq j$, thus satisfying weak positive response. Minimal fairness is also satisfied by noticing the symmetry in the definitions of f_1, f_2.

Now we need to prove that any ranking system F satisfying the four criteria is three-plurality. By non-imposition, there exist graphs G_1, G_2, G_3 such that: $v_0 \prec^F_{G_1} v_1 \prec^F_{G_1} v_2$, $v_2 \prec^F_{G_2} v_0 \prec^F_{G_2} v_1$, and $v_1 \prec^F_{G_3} v_2 \prec^F_{G_3} v_0$. The set of allowable strategies for agent v_i for $i \in \{0, 1, 2\}$ is $\{s^i_1, s^i_2, s^i_3, s^i_4\} = \wp(V \setminus \{v_i\})$. We can use strategy vectors of the form (s^0_i, s^1_j, s^2_k) to represent the graph $(V, s^0_i \cup s^1_j \cup s^2_k)$.

Let $\vec{s_1}, \vec{s_2}, \vec{s_3}$ be the strategy vectors representing G_1, G_2, G_3 respectively. By incentive compatibility, $\vec{s_1}$ and $\vec{s_2}$ differ by the strategies of at least 2 agents. Assume that $s^0_1 \neq s^0_2 \wedge s^1_1 \neq s^1_2 \wedge s^2_1 \neq s^2_2$. By IC, in the graph (s^0_2, s^1_1, s^2_1): $r(v_0) = 0.5$ and in the graph (s^0_2, s^1_1, s^2_2): $r(v_1) = 2.5$. As these two graphs differ only in the outgoing edges of v_2, its rank must be equal, thus must be $r(v_2) = 1.5$ in both. Therefore, in both (s_2^0, s_1^1, s_1^2) and (s_2^0, s_1^1, s_2^2), F ranks $v_0 \prec v_2 \prec v_1$. Again from IC, graph (s_2^0, s_2^1, s_1^2) must be ranked $v_2 \prec v_0 \prec v_1$ and (s_1^0, s_1^1, s_2^2) must be ranked $v_0 \prec v_1 \prec v_2$. We can now let $G_2 = (s_2^0, s_2^1, s_1^2)$, and thus differ from G_1 by the strategies of only two agents.

It is easy to see we can always choose G_2 such that G_1 and G_2 only differ in the strategies of v_0 and v_1. Similarly, G_3 can be chosen such that G_1 and G_3 differ only by the strategies of v_1 and v_2. Assume now that $s_3^1 \neq s_2^1$. By IC, in graph (s_1^0, s_3^1, s_1^2): $r(v_2) = r(v_1) = 1.5$ and thus $v_0 \simeq v_1 \simeq v_2$. Now, in graph (s_1^0, s_2^1, s_1^2): $r(v_0) = r(v_1) = 1.5$ and thus $v_0 \simeq v_1 \simeq v_2$. Now, in graph (s_1^0, s_2^1, s_3^2): $r(v_1) = 0.5$ and $r(v_2) = 1.5$, so $v_1 \prec v_2 \prec v_0$. We now let $G_3 = (s_1^0, s_2^1, s_3^2)$ and thus now every pair of graphs from G_1, G_2, G_3 differ by strategies of two agents. After renaming strategies, we get a structure isomorphic to the one described in Figure 1, but without any mapping between the names of the strategies and acutal edge selection by the agents.

We will first show that the additional strategies of the agents simply reflect these existing strategies. In (s_3^0, s_1^1, s_1^2), by IC, $r(v_0) = 1.5$. So assume that F ranks $v_2 \prec v_0 \prec v_1$. However, in that case in (s_3^0, s_2^1, s_1^2), $r(v_1) = r(v_0) = 2.5$, which is impossible. However, in the two remaining cases it is easy to see that s_3^0 reflects s_1^0 or s_2^0. The same is true for all other agents. Therefore, we only need to map the four strategies for each agent to one of the two options for that agent.

Note that agent v_2 is strengthened when agent v_0 switches from s_1^0 to s_2^0 and agent v_1 is weakened. Assume $S(v_0) = \{v_1\}$ maps to s_2^0, then by weak positive response, $S(v_0) = \{v_1, v_2\}$ and $S(v_0) = \emptyset$ must also map to s_2^0, and furthermore then $S(v_0) = \{v_2\}$ must also map to s_2^0, in contradiction to the fact that s_1^0 must be playable (by non-imposition). Similarly, in all cases where $|S(v)| = 1$, $S(v)$ maps to the relevant strategy in Figure 1.

By minimal fairness, when $E = \emptyset$, the strategy profile must be (s_1^0, s_1^1, s_1^2) or (s_2^0, s_2^1, s_2^2), thus if $S(v_0) = \emptyset$ maps to a strategy s_i^0, then $S(v_1) = \emptyset$ and $S(v_2) = \emptyset$ must map to strategies s_i^1 and s_i^2 respectively. The same goes for $S(v_0) = \{v_1, v_2\}$ — if it maps to a strategy s_i^0, then $S(v_1) = \{v_0, v_2\}$ and $S(v_2) = \{v_0, v_1\}$ must map to strategies s_i^1 and s_i^2 respectively.

So, we are left with four mapping options:

- $S(v_0) = \emptyset$ maps to s_2^0 and $S(v_0) = \{v_1, v_2\}$ maps to s_1^0 .
- $S(v_0) = \emptyset$ maps to s_1^0 and $S(v_0) = \{v_1, v_2\}$ maps to s_2^0 .
- $S(v_0) = \emptyset$ and $S(v_0) = \{v_1, v_2\}$ both map to s_2^0 .
- $S(v_0) = \emptyset$ and $S(v_0) = \{v_1, v_2\}$ both map to s_1^0 .

These mapping options exactly correspond to the four three-plurality ranking systems — The first two correspond to f_1, and the second two to f_2. Of each pair, the first corresponds to the ordering v_0, v_1, v_2 and the second corresponds to v_0, v_2, v_1.

We have shown any ranking system satisfying the four conditions must be three-plurality.

To show that the conditions are independent we must show different ranking systems satisfying all conditions except one:

- Incentive compatibility — The approval voting ranking system satisfies all aforementioned conditions except incentive compatibility.
- Non-imposition — The trivial ranking system that always ranks all vertices equally satisfies IC, weak positive response and minimal fairness.
- Weak positive response — We can swap the meanings of s_1^i and s_2^i for all agents and get a ranking system satisfying all conditions except weak positive response.
- Minimal fairness — If we do not assume minimal fairness, we can assign the strategies for $S(v) = \emptyset$ and $S(v) = V \setminus \{v\}$ differently for each agent v. $\square$

References

Altman, A., and Tennenholtz, M. 2005a. On the axiomatic foundations of ranking systems. In *Proc. 19th International Joint Conference on Artificial Intelligence*, 917–922.

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

Altman, A., and Tennenholtz, M. 2006. Incentive compatible ranking systems. unpublished manuscript, Technion.

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

Borodin, A.; Roberts, G. O.; Rosenthal, J. S.; and Tsaparas, P. 2005. Link analysis ranking: algorithms, theory, and experiments. *ACM Trans. Inter. Tech.* 5(1):231–297.

Cheng, A., and Friedman, E. 2005. Sybilproof reputation mechanisms. In *P2PECON '05: Proceeding of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, 128–132. New York, NY, USA: ACM Press.

Dutta, B.; Jackson, M. O.; and Le Breton, M. 2001. Strategic candidacy and voting procedures. *Econometrica* 69(4):1013–37.

Kleinberg, J. M. 1999. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)* 46(5):604–632.

Mas-Colell, A.; Whinston, M.; and Green, J. 1995. *Microeconomic Theory*. Oxford University Press.

Page, L.; Brin, S.; Motwani, R.; and Winograd, T. 1998. The pagerank citation ranking: Bringing order to the web. Technical Report, Stanford University.

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

Resnick, P., and 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.; and Kuwabara, E. 2000. Reputation systems. *Communications of the ACM* 43(12):45–48.

Slutzki, G., and Volij, O. 2005. Ranking participants in generalized tournaments. *International Journal of Game Theory* 33(2):255–270.

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