An Axiomatic Approach to Personalized Ranking Systems

ALON ALTMAN

Stanford University, Stanford, California

AND

MOSHE TENNENHOLTZ

Technion—Israel Institute of Technology, Haifa, Israel
Microsoft Israel R & D Center

Abstract. Personalized ranking systems and trust systems are an essential tool for collaboration in a multi-agent environment. In these systems, trust relations between many agents are aggregated to produce a personalized trust rating of the agents. In this article, we introduce the first extensive axiomatic study of this setting, and explore a wide array of well-known and new personalized ranking systems. We adapt several axioms (basic criteria) from the literature on global ranking systems to the context of personalized ranking systems, and fully classify the set of systems that satisfy all of these axioms. We further show that all these axioms are necessary for this result.

Categories and Subject Descriptors: G.2.2 [**Discrete Mathematics**]: Graph Theory—*Graph algorithms; social networks*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*; J.4 [**Social and Behavioral Sciences**]: *Economics; sociology*; K.4.4 [**Computers and Society**]: Electronic Commerce—*Ranking systems*

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Axiomatic approach, manipulation, ranking systems, social networks

ACM Reference Format:

Altman, A., and Tennenholtz, M. 2010. An axiomatic approach to personalized ranking systems. J. ACM 57, 4, Article 26, (April 2010), 35 pages.
DOI = 10.1145/1734213.1734220 http://doi.acm.org/10.1145/1734213.1734220

26

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

This article is an extended version of Altman, A., and Tennenholtz, M., An axiomatic approach to personalized ranking systems, in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*.

Authors' addresses: A. Altman, Stanford University, Stanford Computer Science, Gates 258, 353 Serra Mall, Stanford, CA 94305-9020, e-mail: epsalon@stanford.edu; M. Tennenholtz, William Davidson Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, Technion City Haifa 32000, Israel, e-mail: moshet@ie.technion.ac.il.

DOI 10.1145/1734213.1734220 http://doi.acm.org/10.1145/1734213.1734220

1. *Introduction*

Personalized ranking systems and trust systems are an essential tool for collaboration in a multi-agent environment. In these systems, agents report on their peers' performance, and these reports are aggregated to form a ranking of the agents. This ranking may be either global, where a single ranking is generated and reported to all the agents, or personalized, where each agent is provided with her own ranking of the agents. Examples of global ranking systems include eBay's reputation system [Resnick and Zeckhauser 2001] and Google's PageRank [Page et al. 1998]. Examples of personalized ranking systems include personalized versions of PageRank [Haveliwala et al. 2003; Jeh and Widom 2003] and the MoleTrust ranking system [Avesani et al. 2005].

Trust systems which provide each agent with a set of agents he or she can trust can be viewed as personalized ranking systems which supply a two-level ranking over the agents. Many of these systems can be easily adapted to provide a full ranking of the agents. Examples of trust systems include OpenPGP(Pretty Good Privacy)'s trust system [Callas et al. 1998], the ranking system employed by Advogato [Levien 2002], and the epinions.com web of trust.

A central challenge in the study of ranking systems, is to provide means and rigorous tools for the evaluation of these systems. In this article, we focus on the axiomatic approach. In this approach, one considers basic properties, or axioms, one might require a ranking system to satisfy. Then, existing and new systems are classified according to the set of axioms they satisfy. Typical results of such study are axiomatizations of particular ranking systems, or proofs that no ranking system satisfying a set of axioms exists. For example, in Altman and Tennenholtz [2005] we provide a set of axioms that are satisfied by the PageRank system and show that any global ranking system that satisfies these axioms must coincide with PageRank.

While the axiomatic approach has been extensively applied to the global ranking systems setting [Tennenholtz 2004; Altman and Tennenholtz 2005, 2008, 2007], no general attempt has been made to apply such an approach to the context of personalized ranking systems. In this article, we introduce an extensive axiomatic study of the personalized ranking system setting by adapting axioms that have been previously applied to global ranking systems. We compare several existing personalized ranking systems in the light of these axioms, and provide novel ranking systems that satisfy various sets of axioms. Moreover, we prove a full characterization of the personalized ranking systems satisfying all suggested axioms.

We consider four basic axioms. The first axiom, self confidence, requires that an agent would be ranked at the top of his own personalized rank. The second axiom, transitivity, captures the idea that an agent preferred by more highly trusted agents should be ranked higher than an agent preferred by less trusted agents. The third axiom, Ranked Independence of Irrelevant Alternatives, requires that under the perspective of any agent, the relative ranking of two other agents would depend only on the pairwise comparisons between the rank of the agents that trust them. The last axiom, strong incentive compatibility, captures the idea that an agent cannot gain trust in any agent's perspective by manipulating its reported preference or by creating fictitious entities.

We fully characterize the set of ranking systems satisfying all four axioms, and show ranking systems satisfying every three of the four axioms (but not the fourth).

This article is organized as follows: Section 2 introduces the setting of personalized ranking systems and discusses some known systems. In Section 3, we present our axioms, and classify the ranking systems shown according to these axioms. In Section 4, we provide a full characterization of the ranking systems satisfying all of our axioms, and, in Section 5, we study ranking systems satisfying every three of the four axioms. Section 6 presents some concluding remarks and suggestions for future research.

1.1. RELATED WORK. The axiomatic study of personalized ranking systems stems originally from the theory of social choice with Arrow's celebrated Impossibility Theorem [Arrow 1963]. Of specific relevance to ranking systems is the study of voting rules under dichotomous preferences [Bogomolnaia et al. 2005], where agents only have two levels of preference, and the axiomatization of Approval Voting [Fishburn 1978] in this setting.

A central approach to the evaluation of general ranking systems is experimentation. This approach was successfully applied to Hubs&Authorities [Kleinberg 1999] and to various other ranking systems [Borodin et al. 2005]. In the trust systems setting, Massa and Avesani [2005] suggest a similar experimental approach. A major shortcoming of the experimental approach is the lack of a gold standard in order to evaluate the experimental results. The axiomatic approach aims to supply this gold standard.

The axiomatic approach was extensively studied with regard to global ranking systems. Specifically, research has been conducted on the effects of small changes, or perturbations, on ranking systems and the design of systems stable to such changes [Borodin et al. 2005; Ng et al. 2001; Chien et al. 2003; Lee and Borodin 2003; Lempel and Moran 2005]. Axiomatizations were provided for the PageRank ranking system [Altman and Tennenholtz 2005; Palacios-Huerta and Volij 2004], and we have provided some impossibility results in Tennenholtz [2004] and Altman and Tennenholtz [2008].

The issue of incentives in social choice and ranking systems has been also extensively studied from the Gibbard-Satterthwaite theorem [Gibbard 1973; Satterthwaite 1975] for voting to combating link spam on the web [Gyöngyi et al. 2004; Wu et al. 2006; Wu and Davison 2005]. Of particular interest are the recent papers that considered full and almost-full resistance to manipulations by change of outgoing edges [Altman and Tennenholtz 2007, 2006] or creation of fake entities [Cheng and Friedman 2005].

Personalized ranking systems have a lot in common with trust [Dash et al. 2004; Guha et al. 2004] and reputation systems [Resnick et al. 2000; Dellarocas 2003], which try to establish who a specific should agent trust based on social links, similarity, or trade history. These may form the basis for recommender systems which recommend products or services based on similarity [Pennock et al. 2000] or some measure of trust [Andersen et al. 2008].

2. *Personalized Ranking Systems*

2.1. THE SETTING. Before describing our results regarding personalized ranking systems, we must first formally define what we mean by the words "personalized ranking system".

Our definition of a personalized ranking system is ordinal in nature. The resulting personalized ranking only compares vertices but does not allocate numeric values. This practice is common in social choice, as we assume agents are interested in choosing who to interact with and do not care about quantitative values. This can be further motivated by the fact that applications of personalized ranking systems usually require ranked results, with actual values being immaterial (think of search engine results).

Definition 2.1.1. 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 2.1.2. 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 $b \preceq a$.

Personalized ranking systems differ from social welfare functions from the social choice literature in three key aspects:

—The "agents" and "alternatives" coincide—The agents rank themselves. This is the key difference between ranking systems (personalized or otherwise) and social welfare functions and voting rules. This feature emphasizes the fact that agents submit trust votes for each other, and not for some external alternatives.

—A designated source node s is explicitly specified. This aspect is the basis of the personalization. There is no single ranking of the agents as with general ranking systems, but rather a different ranking for each agent.

—The input is binary—Agents only select a preferred set of agents, and do not specify a full ranking. Binary input is required in order to avoid Arrow-style impossibility results [Arrow 1963; Gibbard 1973; Satterthwaite 1975] that arise when agents have more than two levels of ranking. The binary input model is of great importance, as it applies to natural settings such as social networks (you vote for your friends) and web pages (where links are regarded as votes).

With these three aspects in mind, we can now define the notion of a personalized ranking system:

Definition 2.1.3. Let $\mathbb{G}_V$ be the set of all directed graphs $G = (V, E)$ with no parallel edges, but possibly with self-loops.[1] A *personalized ranking system(PRS)* F is a functional that for every finite vertex set V and for every *source* $s \in V$ maps every graph $G \in \mathbb{G}_V$ to an ordering $\preceq^F_{G,s} \in L(V)$.

Notation 2.1.4. We will use $P_G(v)$ and $S_G(v)$ to denote the predecessor set and successor set of v in G, respectively. The subscript G may be omitted when understood from context.

Note that our definition does not assume the existence of a directed path from s to every vertex. However, in some settings, this may be considered a useful

[1] Unless otherwise noted, all our results still apply when self loops are not allowed.

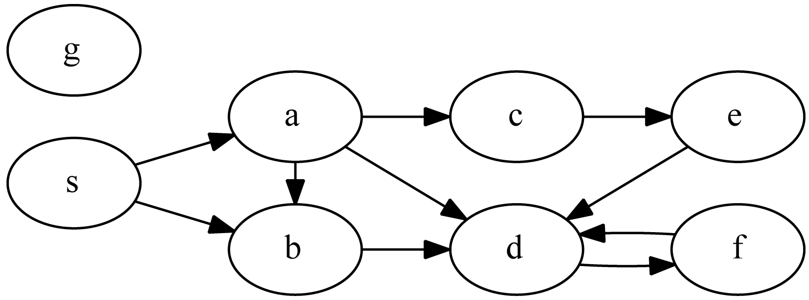

FIG. 1. Example graph for personalized ranking systems.

assumption. Therefore, we shall use these kind of graphs in all examples and counter-examples, but prove our results for the more general case defined above.

2.2. SOME PERSONALIZED RANKING SYSTEMS. We shall now give examples of some known PRSs.

A basic ranking system that is at the basis of many trust systems ranks the agents based on the minimal distance of the agents from the source. The idea behind this system is you trust the most people who you trust directly, and trust declines as farther you go. Many social networking websites implicitly use this type of ranking by having access levels decline with the distance in the friendship graph.

Notation 2.2.1. Let $G = (V, E)$ be some directed graph and $v_1, v_2 \in V$ be some vertices, we will use $d_G(v_1, v_2)$ to denote the length of the shortest directed path in G between v_1 and v_2. If no such path exists, $d_G(v_1, v_2) \triangleq \infty$.

Definition 2.2.2. The *distance PRS F_D* is defined as follows: Given a graph $G = (V, E)$ and a source s, $v_1 \preceq^{F_D}_{G,s} v_2 \Leftrightarrow d_G(s, v_1) \geq d_G(s, v_2)$

Example 2.2.3. Consider the graph in Figure 1. The distance PRS ranks this graph as follows:

$$s \succ a \simeq b \succ c \simeq d \succ e \simeq f \succ g.$$

We see that the rank increases as the distance to s is shorter.

Another family of PRSs can be derived from the well-known PageRank ranking system by modifying the so-called teleportation vector in the definition of PageRank [Jeh and Widom 2003; Haveliwala et al. 2003]. The usual definition of such systems restricts the restart (or “teleportation”) to a personalized set of websites. As in our setting the source of personalization s is represented as a node in the system, we will restrict the restart to only s itself.

The basic idea is that people you should trust are the ones you are most likely to reach by randomly following trust links starting from yourself, balancing between agents who are generally more trusted (and thus more likely to be reached after enough steps) to those which are more trusted by you (and thus reachable soon after the restart).

We begin our formal definition with a classical stochastic matrix defining a random walk on the graph:

Definition 2.2.4. 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:

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

The Personalized PageRank procedure ranks pages according to the stationary probability distribution obtained in the limit of a random walk with a random teleportation to the source s with probability d known as the damping factor. This is formally defined as follows.

Definition 2.2.5. Let $G = (V, E)$ be some graph, and assume $V = \{s, v_2, \ldots, v_n\}$. Let $\mathbf{r}$ be the unique solution of the system

$$(1-d) \cdot \mathbf{r} \cdot A_G + d \cdot (1, 0, \ldots, 0) = \mathbf{r}.$$

The *Personalized PageRank with damping factor d* of a vertex $v_i \in V$ is defined as $PPR^d_{G,s}(v_i) = r_i$, that is, the ith element of the solution $\mathbf{r}$.

The *Personalized PageRank Ranking System with damping factor d* is a PRS that for the vertex set V and source $s \in V$ maps G to $\preceq^{PPR_d}_{G,s}$, where $\preceq^{PPR_d}_{G,s}$ is defined as: for all $v_i, v_j \in V$: $v_i \preceq^{PPR_d}_{G,s} v_j$ if and only if $PPR^d_{G,s}(v_i) \leq PPR^d_{G,s}(v_j)$.

Example 2.2.6. Consider again the graph in Figure 1. The ranking generated by Personalized PageRank depends on the damping factor d being used. Here are the Personalized PageRank values for $d = 0.2$ and $d = 0.5$:

$d =$	s	a	b	c	d	e	f	g
0.2	0.2	0.08	0.1013	0.0213	0.3324	0.0171	0.2579	0
0.5	0.5	0.125	0.14583	0.02083	0.13194	0.01041	0.06597	0

Given these values, the Personalized PageRank ranking system ranks the graph as follows:

$$d = 0.2 \Rightarrow d \succ f \succ s \succ b \succ a \succ c \succ e \succ g$$
$$d = 0.5 \Rightarrow s \succ b \succ d \succ a \succ f \succ c \succ e \succ g.$$

Note that with $d = 0.2$ PPR does not rank s at the top. This seems strange, as you would expect an agent will trust themselves more than any other agent. Indeed, we will later require that s be ranked on top as an axiom.

Further note that both runs of PPR conflict with the distance rule on the ranking of d compared to a—even though a is strictly closer to s than d, PPR ranks d better.

We now suggest a variant of the Personalized PageRank system, which, as we will later show, has more positive properties than Personalized PageRank. The α-Rank system ranks the agents based on their distance from s, breaking ties by summing the values of the predecessors. A small intrinsic value is added to all agents in order to discriminate among agents that have no path from s.

The motivation behind α-rank is to add discriminative power to the distance rule. The tie-breaking scheme used—summation, is based on the intuition from (personalized) PageRank, and as we will see later, ensures that the system satisfies strong transitivity.

Definition 2.2.7. Let $G = (V, E)$ be some graph and assume $V = \{s, v_2, \ldots, v_n\}$. Let B_G be the link matrix for G. That is, $[B_G]_{i,j} = 1 \Leftrightarrow (j, i) \in E$. Let $\alpha = 1/n^2$ and let $\mathbf{a}$ be the unique solution of the system $\alpha \cdot B_G \cdot \mathbf{a} + (1, \alpha^n, \ldots, \alpha^n)^T = \mathbf{a}$. The *$\alpha$-Rank* of a vertex $v_i \in V$ is defined as $r_{G,s}(v_i) = a_i$. The *α-Rank PRS* is a PRS that for the vertex set V and source $s \in V$ maps G to $\preceq^{\alpha R}_{G,s}$, where $\preceq^{\alpha R}_{G,s}$ is defined as: for all $v_i, v_j \in V$: $v_i \preceq^{\alpha R}_{G,s} v_j$ if and only if $r_{G,s}(v_i) \leq r_{G,s}(v_j)$.

By selecting $\alpha = 1/n^2$, it is ensured that a slight difference in rank of nodes closer to s will be more significant than a major difference in rank of nodes further from s, and also ensures the uniqueness of the solution.[2]

We consider α-Rank a variant of Personalized PageRank because both are based on weighted summation of predecessors. The major difference between the two is that in α-Rank there is no scaling based on out-degrees, which we will later see ensures strong transitivity.

Example 2.2.8. Consider again the graph in Figure 1. The α-ranks are as follows:

$$\begin{aligned}
s &= 1\\
a &= \alpha + \alpha^8\\
b &= \alpha + \alpha^2 + \alpha^8 + \alpha^9\\
c &= \alpha^2 + \alpha^8 + \alpha^9\\
e &= \alpha^3 + \alpha^8 + \alpha^9 + \alpha^{10}\\
d &= \frac{2\alpha^2 + \alpha^3 + \alpha^4 + \alpha^8 + 4\alpha^9 + 2\alpha^{10} + \alpha^{11}}{1 - \alpha^2}\\
f &= \frac{2\alpha^3 + \alpha^4 + \alpha^5 + \alpha^8 + \alpha^9 + 2\alpha^{10} + 2\alpha^{11} + \alpha^{12}}{1 - \alpha^2}\\
g &= \alpha^8.
\end{aligned}$$

Now, as α is very small, the α-ranking of this graph must be $s \succ b \succ a \succ d \succ c \succ f \succ e \succ g$.

Additional personalized ranking systems are presented in Section 5 as part of our axiomatic study.

3. *Some Axioms*

In this section, we will present several desirable properties of Personalized Ranking Systems, or axioms. These axioms will form the basis of our study of ranking systems.

3.1. SELF CONFIDENCE. A basic requirement of a PRS is that the source—the agent under whose perspective we define the ranking system—must be ranked strictly at the top of the trust ranking, as each agent implicitly trusts herself. We refer to this property as *self-confidence*.

[2] In the matrix $I - \alpha B_G$, in order for a row to be dependant on other rows some other row must appear with a factor of at least $\frac{1}{\alpha n} = n > 1$, which is impossible.

Definition 3.1.1. Let F be a PRS. We say that F satisfies *self-confidence* if for all graphs $G = (V, E)$, for all sources $s \in V$ and for all vertices $v \in V \setminus \{s\}$: $v \prec^F_{G,s} s$.

3.2. TRANSITIVITY. A basic property of (global) ranking systems called *strong transitivity* [Altman and Tennenholtz 2008; Tennenholtz 2004] requires 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. We adapt this notion to the personalized setting as follows:

Definition 3.2.1. Let F be a PRS. We say that F satisfies *quasi-transitivity* if for all graphs $G = (V, E)$, for all sources $s \in V$ and for all vertices $v_1, v_2 \in V \setminus \{s\}$: Whenever there exists a 1-1 mapping $f : P(v_1) \mapsto P(v_2)$ such that for all $v \in P(v_1)$:$v \preceq f(v)$, then $v_1 \preceq v_2$. F further satisfies *strong quasi-transitivity* if when $P(v_1) \neq \emptyset$ and for all $v \in P(v_1)$: $v \prec f(v)$, then $v_1 \prec v_2$. F further satisfies *strong transitivity* if when either f is not onto or for some $v \in P(v_1)$: $v \prec f(v)$, then $v_1 \prec v_2$.

The definition is based on a matching between the predecessors of v_1 and v_2. This matching ensures that we require v_2 to be at least as strong as v_1 only if the predecessors of v_2 "cover" those of v_1 with at most the same strength.

At this point, transitivity breaks into three gradually stronger variants: The weakest variant, quasi transitivity never induces inequality, and only requires agents with stronger voters to be at least as strong as those with weaker voters.

The second notion, strong quasi transitivity, adds one special case where inequality is required, only if *all* predecessors of one agent map to strictly stronger ones of the other. Note that by itself, strong quasi-transitivity does not induce inequality, as the ranking system $F_=$, which ranks all agents equally does in fact satisfy strong quasi transitivity.

The third and strongest notion, strong transitivity, induces inequality whenever reasonable: When a voters of one agent are stronger than voters of the other, and either one is strictly stronger, or there is additional voter. The latter part of the requirement induces inequality even when this is the only axiom, as it implies, for example, that agents with no voters must be weaker than agents with voters.

Example 3.2.2. A simple example of transitivity is a chain:

$$s \to a \to b$$

If we assume some ranking system F ranks $s \succ a$ (e.g., by self-confidence), then we know that b's predecessor (a) is weaker than a's predecessor (s), and thus quasi transitivity implies $a \succeq b$. If F satisfies strong quasi-transitivity, then we can further infer that $a \succ b$.

Now consider the graph in Figure 1. Any ranking system satisfying quasi transitivity must rank $b \succeq a$, because $P(a) \subseteq P(b)$. Any ranking system satisfying strong transitivity must rank $b \succ a$, because the required mapping from $P(a)$ to $P(b)$ is not onto. Similarly, we can conclude that $e \succ g$ and $d \succ c$, and by further applying strong transitivity also $a \succ c \succ e$ and $f \succ e$.

3.3. RANKED INDEPENDENCE OF IRRELEVANT ALTERNATIVES. A standard assumption in social choice settings is that an alternative'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 the global ranking systems setting [Altman and Tennenholtz 2008], we required that 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. This *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.

We now adapt the axiom of ranked IIA to the setting of PRSs, by requiring this independence for all vertices reachable from the source except the source itself.

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. These possibilities are called comparison profiles:

Definition 3.3.1. A *comparison profile* is a pair $\langle A; B\rangle$ of multisets over $\mathbb{N}$. Let $\mathcal{P}$ denote the set of all such profiles.

A PRS F, a graph $G = (V, E)$, a source $s \in V$, and a pair of vertices $v_1, v_2 \in V$ are said to *satisfy* such a comparison profile $\langle A; B\rangle$ if there exists a function $f : V \mapsto \mathbb{N}$ such that:

$$\begin{aligned}&\forall x, y \in V : x \preceq^F_{G,s} y \Leftrightarrow f(x) \leq f(y),\\&A = f(P(v_1)), \text{ and}\\&B = f(P(v_2)),\end{aligned}$$

where $f(P(v_i))$ is the multiset resulting from the application of f to every predecessor of v_i.

Notation 3.3.2. In order to ease the proofs, we will alternatively consider a comparison profile as a pair of vectors $\langle \mathbf{a}; \mathbf{b}\rangle$ that include the elements of the multisets A and B, respectively, in nondecreasing order.

Comparison profiles reduce an ordinal comparison between the predecessors of two vertices to a pair of multisets of numbers. Each number represents the rank of a predecessor of one of the vertices. For example, the profile $\langle(1, 1); (2)\rangle$ compares two equally ranked weak predecessors (1, 1) to one stronger predecessor (2).

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

Notation 3.3.3. We will use V^G_s to denote the set of vertices that have a directed path from s in a graph G. We will sloppily use V_s when G is understood from context.

Definition 3.3.4. Let F be a PRS. We say that F satisfies *ranked independence of irrelevant alternatives (ranked IIA)* if there exists an ordering $\preccurlyeq$ over multisets over $\mathbb{N}$ such that for every graph $G = (V, E)$, for every source $s \in V$ and for every pair of vertices $v_1, v_2 \in V^G_s \setminus \{s\}$ and for every comparison profile (A, B) that v_1 and v_2 satisfy, $v_1 \preceq^F_{G,s} v_2 \Leftrightarrow A \preccurlyeq B$.

The ranked IIA axiom intuitively means that the relative ranking of agents must be consistent across all comparisons with the same rank relations. For example, there should be one consistent judgement $\prec$ / $\simeq$ / $\succ$ between agents voted by "two weak" (1, 1) compared to those voted by "one strong" (2).

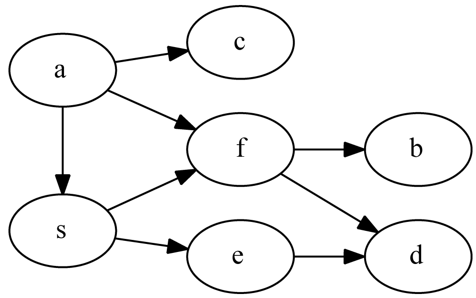

FIG. 2. An example of ranked IIA.

Example 3.3.5. Consider the graph in Figure 2. Furthermore, assume a ranking system F that satisfies ranked IIA has ranked the vertices of this graph as follows:

$$s \succ f \succ e \succ d \succ c \succ b \succ a.$$

Now look at the comparison between e and f. e has one strong predecessor (s) while f has an equal strong predecessor (s again) and a weaker predecessor (a). Therefore, they satisfy the profile $\langle(2);(1,2)\rangle$. Because $e \prec f$, we conclude that $(2) \prec (1,2)$. The same profile occurs again with b (predecessor f) and d (predecessors e and f), and indeed $b \prec d$. If the ranking would have been $d \preceq b$, then we would have known that F does not satisfy ranked IIA.

Now look at b and e. Both have exactly one predecessor, with the former's predecessor weaker than the latter's and thus satisfy the profile $\langle(1);(2)\rangle$. As, $b \prec e$, we conclude that $(1) \prec (2)$. This same profile is satisfied by the pairs (s, e) and (c, b). However, in both cases, we have the opposite result $s \succ e$ and $c \succ b$. These results are still consistent with ranked IIA, because it specifically does not apply to s or to vertices with no directed path from s (such as c).

For F to satisfy ranked IIA, we must show that these and all other comparison profiles maintain the same results here and in all other graphs.

3.4. INCENTIVE COMPATIBILITY. As with global ranking systems, agents ranked by personalized ranking systems may wish to manipulate their reported preferences in order to improve their trustworthiness in the eyes of a specific agent. Therefore, the incentives of these agents should in many cases be taken into consideration.

The issue of incentives has been extensively studied both in classical social choice [Gibbard 1973; Satterthwaite 1975; Dutta et al. 2001], and with regard to global ranking systems [Altman and Tennenholtz 2007, 2006]. In the context of web search, many approaches have been applied to the detection and avoidance of malicious spam websites [Gyöngyi et al. 2004; Wu and Davison 2005; Wu et al. 2006]. These systems usually assume spam pages exist but have a limited number of votes from "good" pages, and try to identify spam with high probability using seeding with trusted sources. Our approach differs from the web-spam approach in that we assume any agent (except the source) can turn malicious and change its outgoing links or create false entities, and the personalized ranking system must inherently deny that agent any benefit from that action, even though it could not be outright detected. Specifically, we require that any user turned malicious cannot improve their position over the one they legitimately have based on their incoming links.

We would like our ranking systems to stand against various types of manipulations. It is important to formally define what a manipulation is, and the types of manipulations we would like to defend against.

Definition 3.4.1. A *manipulation* is a function $\mathcal{M}$ that maps every graph $G = (V, E) \in \mathbb{G}$ and every vertex $v \in V$ in that graph to a set of graphs $M \subseteq \mathbb{G}$ such that $G \in M$ and $v \in G'$ for all $G' \in M$.

We will use manipulations to define for every vertex in any graph, what different graphs can that agent cause to be presented to the ranking system as a result of a manipulation.

3.4.1. *Utility Structure.* We shall now define incentive compatibility, following Altman and Tennenholtz [2007]. We require that a ranking system will not rank agents better when they apply a manipulation, but we assume that the agents are interested only in their own ranking (and not, say, in the ranking of those they prefer).

We assume that for strict rankings (with no ties), there exists a utility function $u : \mathbb{N} \mapsto \mathbb{R}$ that maps an agent's *rank* (i.e., the number of agents ranked above it) to a utility value for being ranked that way. We assume u is nonincreasing; that is, every agent weakly prefers to be ranked higher.

This utility function can be extended to the case of ties, by treating these as a uniform randomization over the matching strict orders. Thus, the utility of an agent with k agents strictly above it and m agents tied is

$$u^*(k, m) = E_{r \in \{k, \ldots, k+m-1\}}[u(r)] = \frac{1}{m} \sum_{i=k}^{k+m-1} u(i),$$

We can now define the utility of a ranking for an agent as follows:

Definition 3.4.2. The *utility* $u_G^F(v)$ of a vertex v in graph $G = (V, E)$ under the ranking system F and utility function u is defined as

$$\begin{aligned} u_G^F(v) &= u^*(|\{v' : v' \succ v\}|, |\{v' : v' \simeq v\}|) \\ &= \frac{1}{|\{v' : v' \simeq v\}|} \sum_{i=|\{v':v' \succ v\}|}^{|\{v':v' \succeq v\}|-1} u(i). \end{aligned}$$

This definition allows us to define a preference relation over rankings for each agent. Using this preference relation, we can now define the general notion of incentive compatibility as immunity to utility increase as a result of a manipulation:

Definition 3.4.3. Let F be a ranking system. F is called *incentive compatible under manipulation $\mathcal{M}$ and utility function u* if for all graphs $G = (V, E)$, for all sources $s \in V$, for all vertices $v \in V$, and for all manipulations $G' \in \mathcal{M}(G, v)$: $u_G^F(v) \geq u_{G'}^F(v)$.

Example 3.4.4. Assume a ranking system F ranks some graph G as follows:

$$s \succ a \succ b \simeq c \simeq d,$$

and assume c can manipulate the graph causing a different graph $G' \in \mathcal{M}(G, c)$ to be ranked by the system. In that case F ranks G' as follows:

$$s \succ a \succ b \succ c \succ d.$$

Now suppose a utility function $u(r) = -r^2$. In G', the utility of c from F is $u^F_{G'}(c) = -3^2 = -9$. In G the utility of c from F is $u^F_G = \frac{u(2)+u(3)+u(4)}{3} = -\frac{4+9+16}{3} = -9.667$. We see that the manipulation by c has increased its utility, and thus F is not incentive compatible under u.

Now suppose a utility function $u'(r) = -r$. In G, the utility of c from F is $u'^F_G(c) = -\frac{2+3+4}{4} = -3$, while in G' it is also $u'^F_{G'}(c) = -3$. Thus, we see that, if this is the only possible manipulation, F is incentive compatible under u'.

One interesting family of utility functions arises if we assume the target agent considers only the first k results returned by the ranking system for some k. For example, this might occur if only k results are presented, or can be read in a given time:

Definition 3.4.5. A utility function $u : \mathbb{N} \mapsto \mathbb{R}$ is termed *threshold* if the exists some $k \in \mathbb{N}$ such that

$$u(i) = \begin{cases} 1 & i < k \\ 0 & \text{Otherwise.} \end{cases}$$

We can now define a strong version of incentive compatibility, requiring incentive compatibility under any (reasonable) utility function.

Definition 3.4.6. Let F be a ranking system. F satisfies *strong incentive compatibility under manipulation* $\mathcal{M}$ if for any nonincreasing utility function $u : \mathbb{N} \mapsto \mathbb{R}$, F is incentive compatible under $\mathcal{M}$ and u.

It turns out that this strong notion of incentive compatibility is entailed from the more limited requirement of incentive compatibility under threshold utility functions.

Furthermore, we show that we can do without utility functions altogether, and define strong incentive compatibility only in the terms of changes to the resulting ranking: For strong incentive compatibility to hold, an agent should be able to improve its position neither by reducing ties, nor by having less agents ranked above it.

LEMMA 3.4.7. *Let F be a PRS, and let $\mathcal{M}$ be a manipulation. Then, the following are equivalent:*

(1) *F satisfies strong incentive compatibility under $\mathcal{M}$.*
(2) *For any threshold utility function u: F is incentive compatible under $\mathcal{M}$ and u.*
(3) *For all graphs $G = (V, E)$, for all sources $s \in V$, for all vertices $v \in V$, and for all manipulations $G' \in \mathcal{M}(G, v)$:*

$$\left|\left\{x \in V' \middle| v \prec^F_{G',s} x\right\}\right| \geq \left|\left\{x \in V \middle| v \prec^F_{G,s} x\right\}\right|, \text{ and}$$
$$\left|\left\{x \in V' \middle| v \preceq^F_{G',s} x\right\}\right| \geq \left|\left\{x \in V \middle| v \preceq^F_{G,s} x\right\}\right|.$$

PROOF. We shall use the following notations throughout the proof:

$$\begin{aligned} n_{\simeq} &= \left|\left\{x \in V \middle| v \simeq^F_{G,s} x\right\}\right| \\ n'_{\simeq} &= \left|\left\{x \in V' \middle| v \simeq^F_{G',s} x\right\}\right| \\ n_{\preceq} &= \left|\left\{x \in V \middle| v \preceq^F_{G,s} x\right\}\right| \\ n'_{\preceq} &= \left|\left\{x \in V' \middle| v \preceq^F_{G',s} x\right\}\right| \end{aligned}$$

(1) $\Rightarrow$ (2) is trivial.

To prove (2) $\Rightarrow$ (3), assume (2) and assume for contradiction that (3) is false. Therefore, there is some graph G, source s, vertex v, and manipulation $G' \in \mathcal{M}(G, v)$ such that either $n'_{\prec} < n_{\prec}$ or $n'_{\preceq} < n'_{\preceq}$. In the former case, the utility function with threshold $n_{\prec}$ will lead to zero utility in G and positive utility in G'. In the latter case the utility function with threshold $n_{\preceq}$ will lead to utility 1 in G' and utility strictly below 1 in G.

To prove (3) $\Rightarrow$ (1), assume (3) and assume for contradiction there is some nonincreasing utility function u, some graph G, source s, vertex v, and manipulation $G' \in \mathcal{M}(G, v)$, such that

$$\frac{1}{n_{\preceq} - n_{\prec}} \sum_{i=n_{\prec}}^{n_{\preceq}-1} u(i) < \frac{1}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i).$$

Now,

$$\begin{aligned} 0 &> \frac{1}{n_{\preceq} - n_{\prec}} \sum_{i=n_{\prec}}^{n_{\preceq}-1} u(i) - \frac{1}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i) \\ &= \frac{\sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i) + \sum_{i=n_{\prec}}^{n'_{\prec}-1} u(i) - \sum_{i=n_{\preceq}}^{n'_{\preceq}-1} u(i)}{n_{\preceq} - n_{\prec}} - \frac{1}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i) \\ 0 &> \frac{n'_{\preceq} - n'_{\prec} - (n_{\preceq} - n_{\prec})}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i) + \sum_{i=n_{\prec}}^{n'_{\prec}-1} u(i) - \sum_{i=n_{\preceq}}^{n'_{\preceq}-1} u(i) \\ &= \left[\frac{n'_{\preceq} - n_{\preceq}}{n'_{\preceq} - n'_{\prec}} - \frac{n'_{\prec} - n_{\prec}}{n'_{\preceq} - n'_{\prec}}\right] \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i) + \sum_{i=n_{\prec}}^{n'_{\prec}-1} u(i) - \sum_{i=n_{\preceq}}^{n'_{\preceq}-1} u(i) \\ &= \frac{n'_{\preceq} - n_{\preceq}}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i) - \sum_{i=n_{\preceq}}^{n'_{\preceq}-1} u(i) + \sum_{i=n_{\prec}}^{n'_{\prec}-1} u(i) - \frac{n'_{\prec} - n_{\prec}}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i) \\ &\geq \frac{n'_{\preceq} - n_{\preceq}}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i) - (n'_{\preceq} - n_{\preceq}) u(n_{\preceq}) \end{aligned}$$

$$+ (n'_{\prec} - n_{\prec})u(n'_{\prec} - 1) - \frac{n'_{\prec} - n_{\prec}}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} u(i)$$

$$= \frac{n'_{\preceq} - n_{\preceq}}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} [u(i) - u(n_{\preceq})] + \frac{n'_{\prec} - n_{\prec}}{n'_{\preceq} - n'_{\prec}} \sum_{i=n'_{\prec}}^{n'_{\preceq}-1} [u(n'_{\prec} - 1) - u(i)] \geq 0,$$

in contradiction to our assumption. □

Even though the notion of strong incentive compatibility stems from utility functions, it is easier to use case (3) of Lemma 3.4.7, and we shall do so throughout the article.

3.4.2. *Manipulations.* In Altman and Tennenholtz [2007, 2006], we considered manipulation by modification of an agent's outgoing links. Such outgoing link manipulation can be defined as:

$$\mathcal{M}_{out}(V, E, v) = \{(V, E') | \forall u \in V \setminus \{v\} : \forall u' \in V : (u, u') \in E \Leftrightarrow (u, u') \in E'\}.$$

Example 3.4.8. Consider the graph in Figure 3(a). If agent a performs an outgoing link manipulation, she can change the graph to include any subset of the dotted links in Figure 3(b). Note that a cannot remove the incoming link from s or add new incoming links from other agents. Furthermore, she cannot add or remove agents in the graph.

The outgoing link manipulation $\mathcal{M}_{out}$ is actually a special kind of manipulation in the sense that the agent can perform the manipulation in both directions.

Definition 3.4.9. A manipulation $\mathcal{M}$ is called *reversible* if for all $G = (V, E) \in \mathbb{G}$, for all $v \in V$, and for all $G' \in \mathcal{M}(G, v)$: $G \in \mathcal{M}(G', v)$.

Reversible manipulations are important due to the following simple fact:

FACT 3.4.10. *Let $\mathcal{M}$ be a reversible manipulation and let F be a PRS . F satisfies strong incentive compatibility under $\mathcal{M}$ if and only if for all graphs $G = (V, E)$, for all sources $s \in V$, for all vertices $v \in V$, and for all manipulations $G' \in \mathcal{M}(G, v)$:*

$$\left|\left\{x \in V' \middle| v \prec^F_{G',s} x\right\}\right| = \left|\left\{x \in V \middle| v \prec^F_{G,s} x\right\}\right|, \textit{and}$$
$$\left|\left\{x \in V' \middle| v \simeq^F_{G',s} x\right\}\right| = \left|\left\{x \in V \middle| v \simeq^F_{G,s} x\right\}\right|.$$

Therefore, in a PRS that is incentive compatible under a reversible manipulation an agent cannot change its rank at all (for better or for worse) by performing a manipulation.

Another type of manipulation, considered by Cheng and Friedman [2005] is concerned with the generation of fraudulent identities in order to manipulate one's rank. Their setting considered weighted edges, as opposed to our setting where the edges are binary. However, we can adapt their sybil form of manipulation by simply removing these weights.

A sybil manipulation, or sybil strategy is a manipulation in which an agent controlling one vertex v in the graph can create any number of fraudulent identities (or sybils) and freely manipulate the links among these sybils, while maintaining

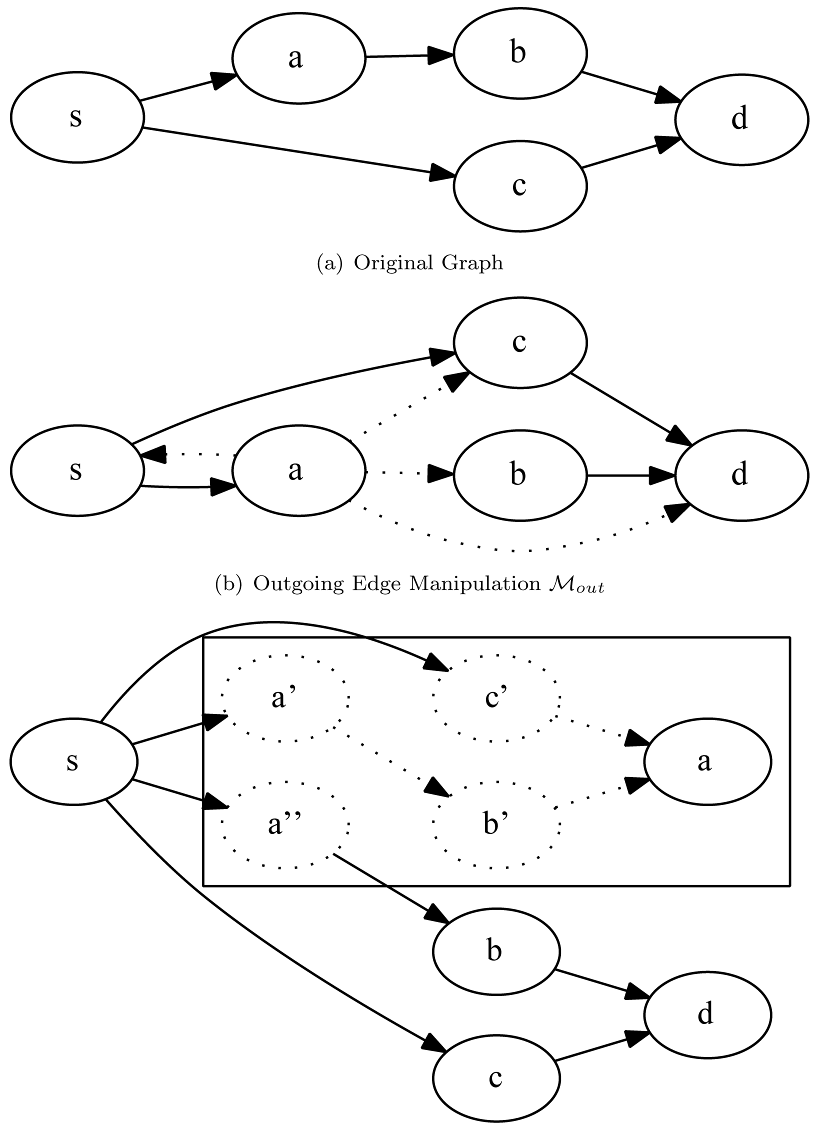

(a) Original Graph

(b) Outgoing Edge Manipulation $\mathcal{M}_{out}$

(c) Sybil Manipulation $\mathcal{M}_{sybil}$

FIG. 3. Examples of manipulations.

the same set of incoming and outgoing links (possibly duplicated) among the sybil group as a whole.

Thus, we can define the sybil manipulation as:

$$\begin{aligned}\mathcal{M}_{sybil}(V, E, v) = \{(V', E') | V' = V \uplus A \\ \wedge \forall u, u' \in V \setminus \{v\} : \\ \wedge (u, u') \in E \Leftrightarrow (u, u') \in E' \\ \wedge (u, v) \in E \Leftrightarrow \exists a \in A \cup \{v\} : (u, a) \in E' \\ (v, u) \in E \Leftrightarrow \exists a \in A \cup \{v\} : (a, u) \in E'.\end{aligned}$$

Example 3.4.11. Consider again the graph in Figure 3(a), and now consider a sybil manipulation by agent a. Now, a can create any number of new vertices and edges between them, with only the limitation of keeping the same incoming and outgoing edges, which may be duplicated.

To illustrate how strong this manipulation is, consider the manipulation in Figure 3(c). Here, a has created a complete replica of the original graph with herself in place of agent d. A ranking system that is not sensitive to node names must rank a and d the same in the manipulated graph. That means that a ranking system that is strongly incentive compatible under $\mathcal{M}_{sybil}$ must assign d in the second graph a rank no better than a in the original graph.

We can also consider the combined manipulation of the two, which is not the same as the simple union of these manipulations:

$$\begin{aligned}\mathcal{M}_{both}(V, E, v) = \{(V', E') | V' = V \uplus A \\ \wedge \forall u, u' \in V \setminus \{v\} : (u, u') \in E \Leftrightarrow (u, u') \in E' \\ \wedge (u, v) \in E \Leftrightarrow \exists a \in A \cup \{v\} : (u, a) \in E'.\end{aligned}$$

It turns out that strong incentive compatibility under both outgoing edge and sybil manipulations is equivalent to strong incentive compatibility under the combined manipulation:

FACT 3.4.12. *Let F be a PRS. F satisfies strong incentive compatibility under $\mathcal{M}_{out}$ and under $\mathcal{M}_{sybil}$ if and only if it satisfies strong incentive compatibility under $\mathcal{M}_{both}$.*

PROOF. The "if" direction is trivial. For the "only if" direction, let $G = (V, E)$ be a graph and $v \in V$. Consider a manipulation $(V', E') \in \mathcal{M}_{both}(V, E, v)$. Let

$$\begin{aligned}U &= \{x | \exists u \in V' \setminus V \cup \{v\} : (u, x) \in E'\} \\ E'' &= E \setminus \{(v, x) | x \in V\} \cup \{(v, x) | x \in U\}.\end{aligned}$$

Now $(V, E'') \in \mathcal{M}_{out}(V, E, v)$ and $(V', E') \in \mathcal{M}_{sybil}(V, E'', v)$, and due to strong incentive compatibility under these manipulations, F also satisfies strong incentive compatibility under manipulation (V', E') and indeed under any manipulation in $\mathcal{M}_{both}$. $\square$

It turns out that for personalized ranking systems, strong incentive compatibility under these two different manipulations almost always goes hand in hand, and for the ranking systems, we present in this article we will see that they either satisfy strong incentive compatibility under $\mathcal{M}_{both}$ or do not satisfy incentive compatibility under neither $\mathcal{M}_{out}$ nor $\mathcal{M}_{sybil}$.

3.5. SATISFICATION. We will now demonstrate the aforementioned axioms by showing which axioms are satisfied by the PRSs mentioned in Section 2.2.

PROPOSITION 3.5.1. *The distance PRS F_D satisfies self confidence, ranked IIA, transitivity, and strong incentive compatibility under $\mathcal{M}_{both}$, but does not satisfy strong transitivity.*

PROOF. Self-confidence is satisfied by definition of F_D. F_D satisfies ranked IIA, because it ranks every comparison profile in the connected section consistently

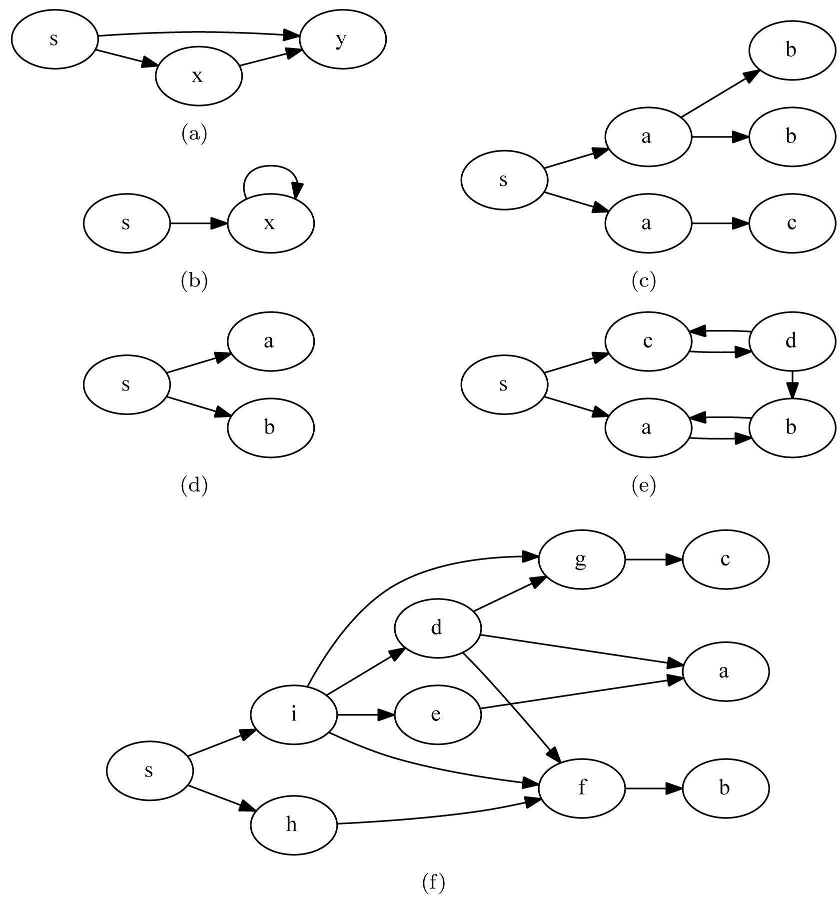

FIG. 4. Graphs proving PRS do not satisfy axioms.

according to the following rule:

$$(a_1, a_2, \ldots, a_n) \preccurlyeq (b_1, b_2, \ldots, b_m) \Leftrightarrow a_n \leq b_m.$$

That is, any two vertices are compared according to their strongest predecessor. F_D satisfies strong quasi transitivity, because the ranking of the profiles above is consistent with strong quasi transitivity. The unconnected vertices are all equal to each other and weaker than the connected vertices which is also true for their predecessors, and thus strong quasi transitivity is satisfied.

To prove that F_D satisfies strong incentive compatibility, note the fact that an agent x cannot modify the shortest path from s to x by changing its outgoing links or adding sybils since any such shortest path necessarily does not include x or its sybils (except as target). Moreover, x or its sybils cannot change the shortest path to any agent y with $d(s, y) \leq d(s, x)$, because x and its sybils are necessarily not on the shortest path from s to y. Therefore, the amount of agents ranked above x and its sybils and the amount of agents ranked equal to x or its sybils cannot decrease due to x's manipulations.

To prove F_D does not satisfy strong transitivity, consider the graph in Figure 4(a). In this graph, x and y are ranked the same, even though $P(x) \subsetneq P(y)$, in contradiction to strong transitivity. □

Even though the distance PRS satisfies all our axioms for personalized ranking systems, it suffers from a major problem: On a typical graph (say, for a social network) the second and third levels for any particular source will include a very large number of individuals. Therefore, the generated ranking will not be very useful for the agent.

That said, our result above hints that the distance rule is a good first step towards the design of personalized ranking systems. In particular, we will later see examples of ranking systems that refine the distance rule to generate a ranking with more distinctive power than plain distance.

PROPOSITION 3.5.2. *The Personalized PageRank ranking systems satisfy self confidence if and only if the damping factor is set to more than* $\frac{1}{2}$.[3] *Moreover, Personalized PageRank does not satisfy weak transitivity, ranked IIA or strong incentive compatibility under* $\mathcal{M}_{out}$ *or* $\mathcal{M}_{sybil}$ *for any damping factor.*

PROOF. To prove the that PPR does not satisfy self-confidence for $d \leq \frac{1}{2}$, consider the graph in Figure 4(b). For any damping factor d, the PPR will be $PPR(s) = d$ and $PPR(x) = 1 - d$. If $d \leq \frac{1}{2}$, then $PPR(s) \leq PPR(x)$ and thus $s \preceq^{PPR_d} x$, in contradiction to the self confidence axiom.

PPR satisfies self-confidence for $d > \frac{1}{2}$ because then $PPR(s) \geq d > \frac{1}{2}$, while for all $v \in V \setminus \{s\}$, $PPR(v) \leq 1 - d < \frac{1}{2}$.

To prove that PPR does not satisfy strong quasi transitivity and ranked IIA, consider the graph in Figure 4(c). The PPR of this graph for any damping factor d is as follows: $PPR(s) = d$; $PPR(a) = \frac{d(1-d)}{2}$; $PPR(b) = \frac{d(1-d)^2}{4}$; $PPR(c) = \frac{d(1-d)^2}{2}$. Therefore, the ranking of this graph is: $b \prec c \prec a \prec s$. Quasi-transitivity is violated because $b \prec c$ even though $P(b) = P(c) = a$. This also violates ranked IIA because the ranking profile $\langle(1);(1)\rangle$ must be ranked as equal due to trivial comparisons such as a and a.

Strong incentive compatibility under $\mathcal{M}_{out}$ is not satisfied, because in the graph in Figure 4(c), if any of the b agents b' would have voted for themselves, they would have been ranked $b \prec b' \prec c \prec a \prec s$, which is a strict increase in b' rank.

To show that strong incentive compatibility under $\mathcal{M}_{sybil}$ is not satisfied, consider the graph in Figure 4(d). Note that $a \simeq b \prec s$ in this graph. Consider the manipulation by a where a sybil a' is added along with the edges $\{(s, a'), (a', a)\}$. In this case, the PageRank value of b would be $\frac{1}{3}(1-d)d$ while the PageRank value of a will be $\frac{(1-d)+1}{3}(1-d)d$. Therefore, $b \prec a \prec s$ in the manipulated graph, and thus strong incentive compatibility is not satisfied. □

We have shown that Personalized PageRank is not sybil-proof. In fact, this result is true even if we ignore the effect of the sybils dividing the "strength" of the preceding node. Even in the weighted case, as considered by Cheng and Friedman [2005], replacing a node with a clique may decrease the weight of a single outgoing link and thus increase the relative rank of a desired node.

An example of a PRS that satisfies Strong transitivity is α-Rank:

PROPOSITION 3.5.3. *The* α*-Rank system satisfies self confidence and strong transitivity, but does not satisfy ranked IIA or strong incentive compatibility under* $\mathcal{M}_{out}$ *or* $\mathcal{M}_{sybil}$.

[3] If we do not allow self-loops, this bound becomes $(\sqrt{5} - 1)/2 \approx 0.618$.

PROOF. To show α-Rank satisfies self confidence, note that by definition $r_{G,s}(s) \geq 1$. Assume for contradiction that $\max_{v \neq s} r_{G,s}(v) \geq 1$. Then,

$$\begin{aligned}
r_{G,s}(s) &\leq 1 + \alpha \sum_{v \in V} r_{G,s}(v) \\
&\leq 1 + \alpha \Big[(n-1) \max_{v \neq s} r_{G,s}(v) + r_{G,s}(s) \Big] \\
r_{G,s}(s) &\leq \frac{1}{1-\alpha} + \frac{\alpha}{1-\alpha}(n-1) \max_{v \neq s} r_{G,s}(v) \\
&\leq \frac{4}{3} + \max_{v \neq s} r_{G,s}(v) \\
\max_{v \neq s} r_{G,s}(v) &\leq \alpha^n + \alpha \sum_{v \in V} r_{G,s}(v) \\
&\leq \alpha^n + \alpha \Big[n \cdot \max_{v \neq s} r_{G,s}(v) + \frac{4}{3} \Big] \\
\Big[1 - \frac{n}{n^2} \Big] \max_{v \neq s} r_{G,s}(v) &\leq \frac{4}{3n^2} + \frac{1}{n^{2n}} < \frac{2}{n^2} \\
n^2 - n &< 2 \\
2 \leq n(n-1) &< 2.
\end{aligned}$$

To prove α-Rank satisfies strong transitivity, consider two vertices $a, b \in V \setminus \{s\}$ and a function $f : P(a) \mapsto P(b)$ such that $v \preceq f(v)$ for all $v \in P(a)$. Then,

$$\begin{aligned}
r_{G,s}(a)/\alpha - \alpha^{n-1} &= \sum_{v \in P(a)} r_{G,s}(v) \leq \sum_{v \in f(P(a))} r_{G,s}(v) \\
&\leq \sum_{v \in P(b)} r_{G,s}(v) = r_{G,s}(b)/\alpha - \alpha^{n-1}, \qquad (1)
\end{aligned}$$

which implies $a \preceq b$. If, for some $v \in P(a)$: $v \prec f(v)$, or if f is not onto, then the first or the second inequality, respectively, in (1) above is strict, which implies $a \prec b$, as required.

To prove α-Rank does not satisfy strong incentive compatibility under $\mathcal{M}_{out}$, consider the graph in Figure 4(e). In this graph, α-Rank ranks $d \prec b$. However, if d removes the link to b, they will be ranked equally and thus reducing the number of agents stronger than d. To prove α-Rank does not satisfy strong incentive compatibility under $\mathcal{M}_{sybil}$, consider again the graph in Figure 4(e). Agent c is ranked below agent b in this graph. However, she can duplicate herself and add edges (c, c') and (c', c) to be ranked above b thus decreasing the number of agents ranked better than herself.

To prove α-Rank does not satisfy ranked IIA, consider the graph in Figure 4(f). It is easy to calculate the following α-Rank values:

$$\begin{aligned}
r(s) &= 1 \\
r(i) = r(h) &= \alpha + \alpha^{10} \\
r(d) = r(e) &= \alpha^2 + \alpha^{10} + \alpha^{11} \\
r(f) &= 2\alpha^2 + \alpha^3 + \alpha^{10} + 3\alpha^{11} + \alpha^{12}
\end{aligned}$$

$$
\begin{aligned}
r(g) &= \alpha^2 + \alpha^3 + \alpha^{10} + 2\alpha^{11} + \alpha^{12} \\
r(a) &= 2\alpha^3 + \alpha^{10} + 2\alpha^{11} + 2\alpha^{12} \\
r(b) &= 2\alpha^3 + \alpha^4 + \alpha^{10} + \alpha^{11} + 3\alpha^{12} + \alpha^{13} \\
r(c) &= \alpha^3 + \alpha^4 + \alpha^{10} + \alpha^{11} + 2\alpha^{12} + \alpha^{13}.
\end{aligned}
$$

Therefore, this graph is ranked $c \prec a \prec b \prec d \simeq e \prec g \prec f \prec i \simeq h \prec s$. Note that (a, b) and (a, c) both satisfy the profile $\langle(1, 1); (2)\rangle$, however $a \prec b$ and $c \prec a$ in contradiction to ranked IIA. □

The α-rank system, being additive in nature, satisfies strong transitivity. The small values assigned to all vertices are to ensure a positive value for each vertex, and thus conform to the strong requirement of transitivity. The small value used for α is in order to ensure self-confidence, so that every level will be ranked strictly weaker than the preceding one.

4. *A Characterization Theorem*

We have previously seen that the distance ranking system satisfies all the aforementioned axioms, except strong transitivity, which is not compatible with ranked IIA.[4] However, we have seen that the distance rule does not have much discriminatory power. Therefore, we would like to consider a refinement of the distance rule that would let us discriminate between the vertices that would have otherwise been all ranked equally.

In this section, we suggest to refine the distance rule based on the strength and *number* of strong votes from the previous level. We show that the ranking systems that satisfy these axioms are exactly the ones on the spectrum between the distance rule and this count-based refinement.

We call the systems on this spectrum *strong count* systems. Strong count systems rank agents first according to distance from s, then by their strongest predecessors, breaking ties according to a function of the *number* of equal strongest predecessors the agents have.

The strong count systems on this spectrum differ by a tie-breaking function r. Instead of comparing breaking ties based on the actual predecessor count x, $r(x)$ is used, which may artificially create ties. In the most extreme cases, we get the distance rule for $r \equiv 1$, and the most refined strong count system (as described above) for the identity function $r(x) = x$. For a detailed example, see Example 4.2.

The strong count system is formally defined as follows:

Definition 4.1. Let $r : \mathbb{N} \mapsto \mathbb{N}$ be a monotone nondecreasing function such that $r(i) \leq i$ for all $i \in \mathbb{N} \cup \{0\}$. The strong count system SC_r is defined as follows:

Consider the layers of the graph G based on distance from s:

$$
\begin{aligned}
L_i &= \{x | d(s, x) = i\} \quad \forall i \in \mathbb{N} \cup \{0\} \\
L_\infty &= V \setminus V_s.
\end{aligned}
$$

[4] This has been proven in the global ranking systems setting [Altman and Tennenholtz 2008], and that result copies to the personalized setting.

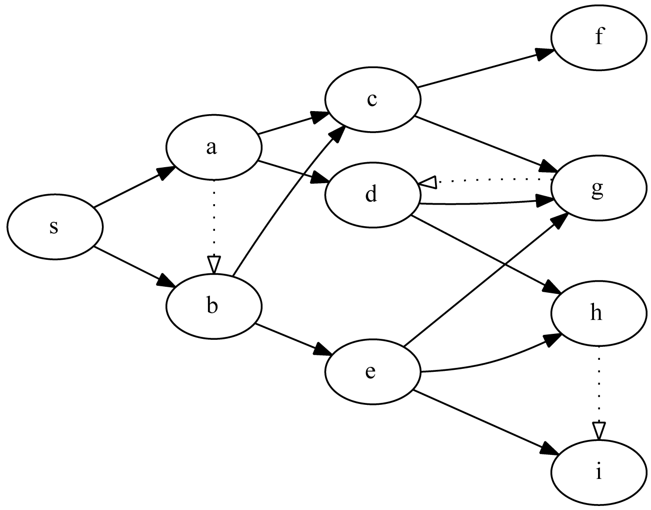

FIG. 5. Example of the Strong Count personalized ranking systems.

Note that all layers are disjoint and are a partition of the vertices:

$$V = \biguplus_{i \in \mathbb{N} \cup \{0,\infty\}} L_i.$$

We shall now define recursively define the order $\preceq_i$ over each layer L_i: Let $M_i(x)$ be a maximum predecessor of x in layer i and $m_i(x)$ be the number of such predecessors:

$$\begin{aligned} M_i(x) &= \max_{\preceq_i}(P(x) \cap L_i) \\ m_i(x) &= |\{v : v \in P(x) \cap L_i \wedge v \simeq_i M_i(x)\}|. \end{aligned} \tag{2}$$

The ordering in layer i is first on the strongest predecessor from layer $i-1$ and second on the number of such predecessors:

$$\begin{aligned} x \preceq_i y \Leftrightarrow\ & [M_{i-1}(x) \prec_{i-1} M_{i-1}(y)] \\ & \vee [M_{i-1}(x) \simeq_{i-1} M_{i-1}(y) \wedge r(m_{i-1}(x)) \leq r(m_{i-1}(y))] \end{aligned}$$

In L_∞ all agents are considered equal:

$$\forall x, y \in L_\infty : x \simeq_\infty y$$

Now,

$$\begin{aligned} a \preceq^{SC_r} b \Leftrightarrow\ & [d(s,a) > d(s,b)] \\ & \vee \left[d(s,a) = d(s,b) \wedge a \preceq_{d(s,a)} b\right]. \end{aligned}$$

Example 4.2. Consider for example the graph in Figure 5. All strong count systems ignore the dotted edges, because these edges are not from a level closer to s, and thus are never the strongest predecessors.

The distance rule, or strong count with $r \equiv 1$, ranks the graph purely based on distance from s:

$$s \succ a \simeq b \succ c \simeq d \simeq e \succ f \simeq g \simeq h \simeq i.$$

If we use strong count with the identity function $r(x) = x$, then a distinction is made in the second level with c having two strong predecessors. This distinction is then carried over to the next level with f and g being stronger than h and i. Between h and i, h is ranked higher due to the the two strong links from d and e, while f and g are equally ranked because they both have one strong predecessor—c. Therefore, the resulting ranking is

$$s \succ a \simeq b \succ c \succ d \simeq e \succ f \simeq g \succ h \succ i.$$

To demonstrate the effect of r, consider now a case where $r(1) = r(2) = 1$ while $r(3) = 2$. In this case, c and d are ranked equally because r has no distinction between two and one strong predecessors. For the same reason f, h, and i are ranked the same. However, as g has three strong predecessors, it is still ranked higher than f, resulting in the ranking

$$s \succ a \simeq b \succ c \simeq d \simeq e \succ g \succ f \simeq h \simeq i.$$

Now we can state our main theorem:

THEOREM 4.3. *Let F be a PRS. The following three statements are equivalent:*

(1) *F is a strong count system for some function r.*
(2) *F satisfies self confidence, strong quasi transitivity, ranked IIA and strong incentive compatibility under $\mathcal{M}_{out}$.*
(3) *F satisfies self confidence, strong quasi transitivity, ranked IIA and strong incentive compatibility under $\mathcal{M}_{both}$.*

We begin our proof by showing that the strong count systems do in fact satisfy all these axioms.

PROOF ($1 \Rightarrow 3$). Let r be a monotone nondecreasing function such that $r(x) \leq x$. SC_r satisfies self-confidence by definition.

To show that SC_r satisfies ranked IIA and strong quasi-transitivity on elements of V_s, we will show that it ranks any profile $p = \langle (a_1, \ldots, a_n); (b_1, \ldots, b_m) \rangle$ as follows: Let $c_a = \max\{i \in \mathbb{N} | a_{n-i} = a_{n-i+1} = \cdots = a_n\}$ and $c_b = \max\{i \in \mathbb{N} | b_{m-i} = b_{m-i+1} = \cdots = b_m\}$.

$$\begin{aligned} \mathbf{a} \preccurlyeq \mathbf{b} \Leftrightarrow\ & (a_n < b_m) \\ & \vee[(a_n = b_m) \wedge (r(c_a) \leq r(c_b))] \end{aligned}$$

This almost follows from the definition of SC_r; however, it remains to show that for all $x \in L_i$: $P(x) \cap L_{i-1} \neq \emptyset$ and for all $z \in P(x) \cap L_{i-1}$, $y \in P(x) \setminus L_{i-1}$: $y \prec z$, thus showing that the limited maximum computed in (2) is equal to the maximum over all $P(x)$. This statement is true due to the strict ordering between layers and the fact that $P(x)$ includes only elements of layer $i-1$ or and layers farther away from s.

Strong quasi-transitivity involving elements in $V \setminus V_s$ and elements either in $V \setminus \{s\}$ is satisfied because for all $x \in V \setminus V_s$ and $y \in V \setminus \{s\}$ we have $x \preceq y$ (by definition) and if $x \prec y$ then $y \in V_s \setminus \{s\}$ and thus there is some $y' \in P(y)$ such that for all $x' \in P(x)$: $x' \preceq y'$.

With regard to the strong incentive compatibility under $\mathcal{M}_{both}$, due to fact the ranking is a refinement of the distance rule, all sybils of v will be strictly weaker than the vertices with smaller distance from s. Furthermore, any other vertices

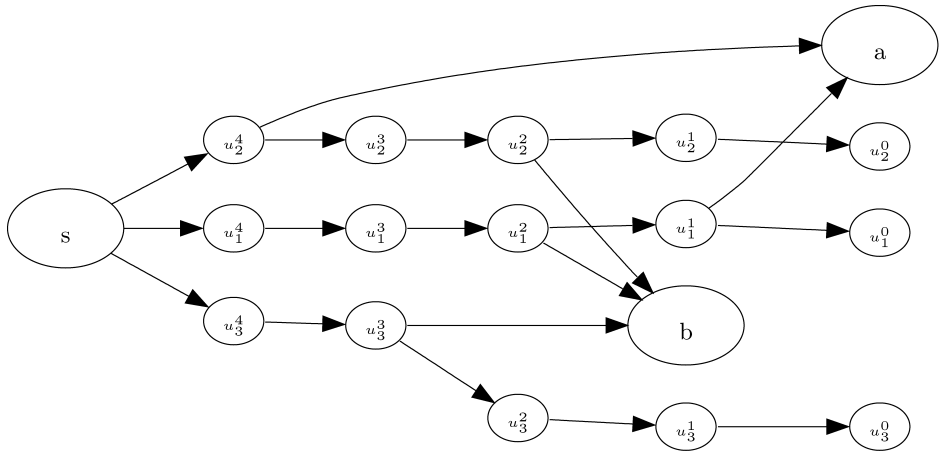

FIG. 6. Example of graph from proof of Lemma 4.5.

that were stronger than v in the original graph will be stronger than any of v's sybils, due to the fact that the relative rank of two vertices is determined only based on incoming links from vertices closer to s, and more incoming edges cannot decrease an agent's rank. By the same logic, vertices which were equal to v in the original graph, will either be stronger or equal to v in the manipulated graph. □

In order to prove the hard direction of Theorem 4.3 ($2 \Rightarrow 1$), we will first show that a strong notion of transitivity is implied by the axioms:

Definition 4.4. Let F be a PRS. We say that F satisfies *weak maximum transitivity* if for all graphs $G = (V, E)$, for all sources $s \in V$ and for all vertices $v_1, v_2 \in V_s$: Let m_1, m_2 be the maximally ranked vertices in $P(v_1)$, $P(v_2)$ respectively. If $m_1 \prec m_2$, then $v_1 \prec v_2$.

LEMMA 4.5. *Let F be a PRS that satisfies self-confidence, strong quasi-transitivity, ranked IIA and strong incentive compatibility. Then, F satisfies weak maximum transitivity.*

PROOF. In order to show that F satisfies weak maximum transitivity, we will show that for every comparison profile the ranking must be consistent with weak maximum transitivity. Let $p = \langle(a_1, a_2, \ldots, a_k); (b_1, b_2, \ldots, b_l)\rangle$ be a comparison profile where $a_k \neq b_l$. Assume without loss of generality that $b_l < a_k$ and assume for contradiction that $(a_1, a_2, \ldots, a_k) \preccurlyeq (b_1, b_2, \ldots, b_l)$. Consider the graph $G = (V, E)$ defined as follows:

$$\begin{aligned} V &= \{s, a, b\} \cup \left\{u_i^j | i \in \{1, \ldots, \max(k, l)\right\}; j \in \{0, \ldots, a_k\}\} \\ E &= \left\{\left(u_i^j, u_i^{j-1}\right) | i \in \{1, \ldots, \max(k, l)\right\}; j \in \{1, \ldots, a_k\}\} \\ &\quad \cup \left\{\left(s, u_i^{b_l}\right) \middle| i \in \{1, \ldots, \max(k, l)\}\right\} \\ &\quad \cup \left\{\left(u_i^j, a\right) \middle| a_i = j\right\} \cup \left\{(u_i^j, b) | b_i = j\right\}. \end{aligned}$$

Figure 6 contains such a graph for the profile $\langle(1, 4); (2, 2, 3)\rangle$.

Note that by strong quasi-transitivity and self confidence, for all i, i', j, j': $u_i^j \preceq u_{i'}^{j'}$ iff $j \leq j'$. Therefore, we will use u^j to denote any u_i^j. By the construction of G, a and b satisfy p. Thus, from our assumption, $a \preceq b$.

By strong quasi-transitivity, $a \succeq u^{b_l}$, and thus from our assumption also $b \succeq u^{b_l}$. Now consider the point of view of agent $u_l^{b_l}$. She can perform a manipulation by not voting for b. This manipulation must not change her relative rank, as it is in $\mathcal{M}_{out}$. As the relative ranks of the u_i^j agents and s are unaffected by this manipulation, it cannot affect the ranks of a and b relative to $u_l^{b_l}$, and thus after the edge $(u_l^{b_l}, b)$ is removed, we still have $b \succeq u_l^{b_l}$. We can repeat this process for all $i = b_l, \ldots, 2$, with the result that in the graph G' for the profile $\langle(a_1, a_2, \ldots, a_k); (b_1)\rangle$, $b \succeq u^{b_2} \succeq u^{b_1}$. However, by strong quasi transitivity, $b \simeq_{G'} u^{b_1-1} \prec_{G'} u^{b_1} \preceq_{G'} b$, which is a contradiction. $\square$

We can now prove the hard direction of Theorem 4.3.

PROOF OF THEOREM 4.3 ($2 \Rightarrow 1$). Given Lemma 4.5, in order to determine the ranking of vertices in V_s it remains to look at profiles $\langle(a_1, a_2, \ldots, a_k); (b_1, b_2, \ldots, b_l)\rangle$ where $a_k = b_l$. Denote $M = a_k = b_l$. Let p be such a profile. Denote $x_a = |\{n|a_n = M\}|$ and similarly $x_b = |\{n|b_n = M\}|$. These values denote the number of strongest predecessors a and b have in profile p.

We will now prove by induction on $k + l - x_a - x_b$ that F ranks p the same as it ranks $\langle(\underbrace{1, \ldots, 1}_{x_a \text{ times}}); (\underbrace{1, \ldots, 1}_{x_b \text{ times}})\rangle$. If $k + l - x_a - x_b = 0$, then $a_1 = a_k = b_1 = b_l$, and thus the requirement is trivially satisfied. Otherwise, we assume correctness for $k + l - x_a - x_b - 1$. Further assume without loss of generality that $a_1 \neq a_k$. Denote $r = a_{k-x_a}$ and $y_a = |\{n|a_n = r\}|$.

We shall now consider two cases:

—If $b_1 = b_l$ or $a_{k-x_a} \neq b_{l-x_b}$. If $b_1 \neq b_l$, then further assume without loss of generality that $a_{k-x_a} > b_{l-x_b}$. Consider the graph $G = (V, E)$ defined as follows:

$$\begin{aligned} V &= \{s, a\} \cup \{b^1, \ldots, b^{y_a}\} \\ &\quad \cup \left\{u_i^j | i \in \{1, \ldots, \max(k, l)\}; j \in \{0, \ldots, M\}\right\} \\ E &= \left\{\left(u_i^j, u_i^{j-1}\right) | i \in \{1, \ldots, \max(k, l)\}; j \in \{1, \ldots, M\}\right\} \\ &\quad \cup \left\{\left(s, u_i^M\right) | i \in \{1, \ldots, \max(k, l)\}\right\} \cup \left\{(u_i^j, a) | a_i = j \neq r\right\} \\ &\quad \cup \left\{\left(u_i^j, b^n\right) | b_i = j, n = 1, \ldots, y_a\right\} \cup \{(b^n, a) | n = 1, \ldots, y_a\}. \end{aligned}$$

Figure 7 contains such a graph for the profile $\langle(1, 3, 3, 4); (1, 2, 4, 4)\rangle$. Note that by strong quasi-transitivity and self-confidence, for all i, i', j, j': $u_i^j \preceq u_{i'}^{j'}$ iff $j \leq j'$. Therefore, we will use u^j to denote any u_i^j. Similarly, all b^n are equal to each other, and by weak maximum transitivity (Lemma 4.5), $u^{M-1} \preceq a, b \prec u^M$ (we will similarly use b to denote any b^n). Therefore, a and b satisfy p. Now consider the following manipulation by b^1: Removing the outgoing edge to a. This manipulation is in $\mathcal{M}_{out}$ and thus should not change the relative rank of b^1. Note that b^1's predecessors remain the same and equal to the ones of $b^2, \ldots, b^{y_a}$, and all b^n remain equal. We must now show that for every allowable relative

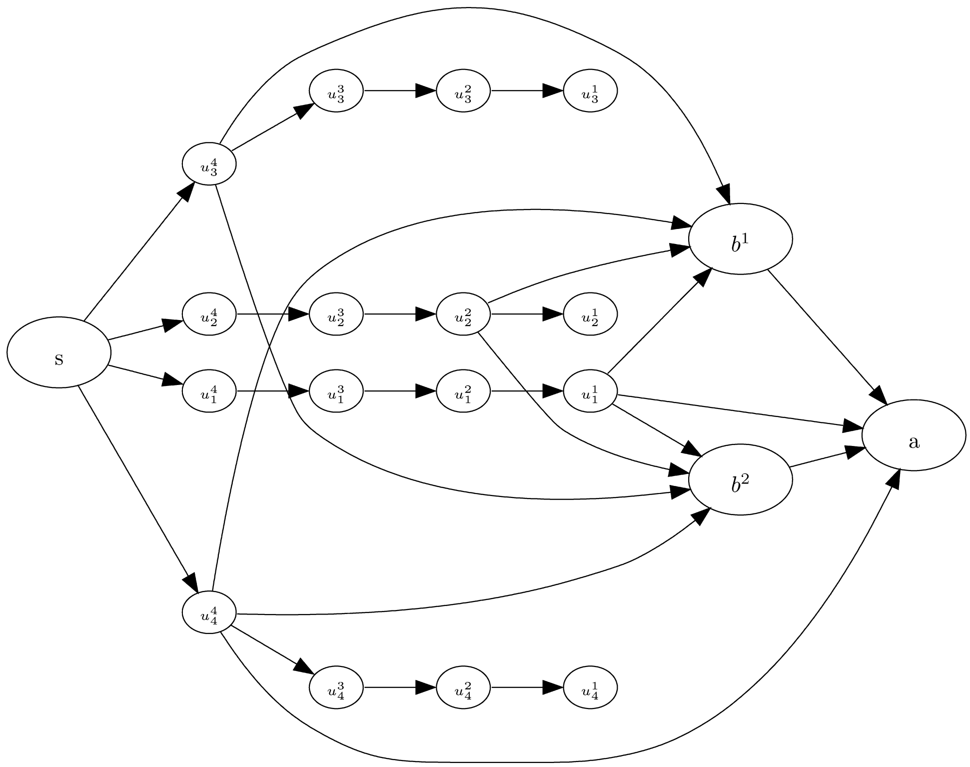

FIG. 7. Example graph from the proof of Theorem 4.3 case 1.

ranking of u^{M-1}, a, and b the manipulation cannot change a and b's relative rank. We will do this by considering all cases:

Ordering	# Vertices equal to b	# Vertices stronger than b
$u^{M-1} \simeq b \prec a$	$y_a + \max(k, l)$	$(M - r) \cdot \max(k, l) + 2$
$u^{M-1} \prec b \prec a$	y_a	$(M - r) \cdot \max(k, l) + 2$
$u^{M-1} \simeq a \simeq b$	$y_a + \max(k, l) + 1$	$(M - r) \cdot \max(k, l) + 1$
$u^{M-1} \prec a \simeq b$	$y_a + 1$	$(M - r) \cdot \max(k, l) + 1$
$u^{M-1} \simeq a \prec b$	y_a	$(M - r) \cdot \max(k, l) + 1$
$u^{M-1} \prec a \prec b$	y_a	$(M - r) \cdot \max(k, l) + 1$

We see that any change in the relation between a and b will surely change b's rank in a way that is not strategy-proof.

We have shown that profile p must be ranked the same as the profile

$$\langle (a_1, a_2, \ldots, a_{k-x_a-1}, a_{k-x_a+1}, \ldots, a_k); (b_1, b_2, \ldots, b_l) \rangle,$$

which by the assumption of induction gives us the desired result.

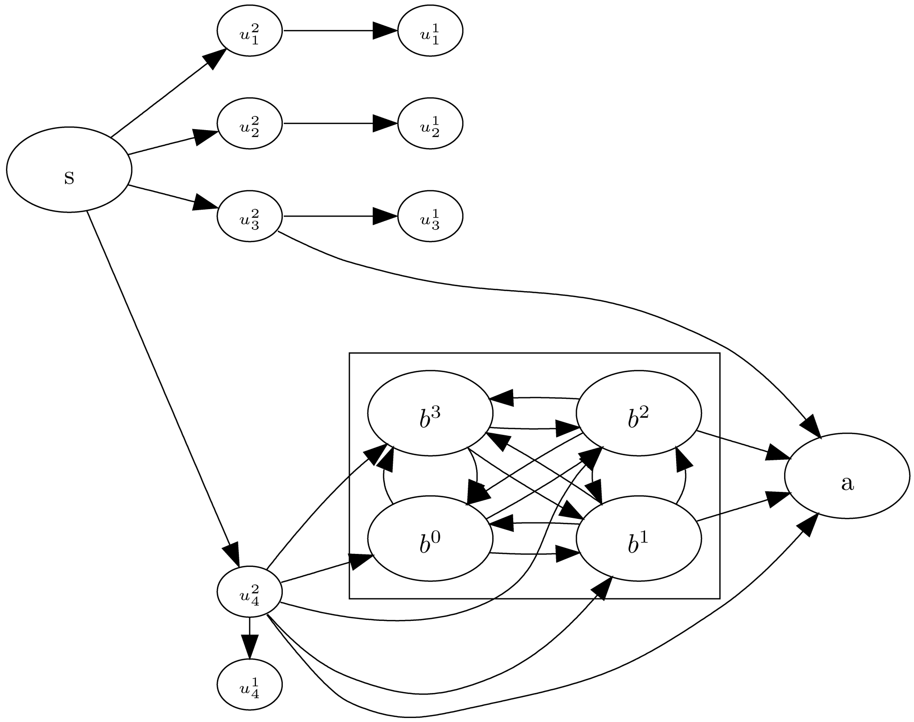

FIG. 8. Example graph from the proof of Theorem 4.3 case 2.

—Otherwise, $a_{k-x_a} = b_{l-x_b}$. Denote $y_b = |\{n|b_n = r\}|$ and assume without loss of generality that $y_b \geq y_a$. Consider the graph $G = (V, E)$ defined as follows:

$$\begin{aligned} V &= \{s, a\} \cup \{b^0, \ldots, b^{y_b}\} \\ &\quad \cup \left\{u_i^j | i \in \{1, \ldots, \max(k, l)\}; j \in \{0, \ldots, M\}\right\} \\ E &= \left\{\left(u_i^j, u_i^{j-1}\right) | i \in \{1, \ldots, \max(k, l)\}; j \in \{1, \ldots, M\}\right\} \\ &\quad \cup \left\{\left(s, u_i^M\right) | i \in \{1, \ldots, \max(k, l)\}\right\} \cup \left\{(u_i^j, a) | a_i = j \neq r\right\} \\ &\quad \cup \left\{\left(u_i^j, b^n\right) | b_i = j \neq r, n = 0, \ldots, y\right\} \\ &\quad \cup \left\{\left(b^n, a\right) | n = 1, \ldots, y_a\right\} \cup \{(b^n, b^m) | n \neq m \in \{0, \ldots, y_b\}\}. \end{aligned}$$

Figure 8 contains such a graph for the profile $\langle(1, 1, 2, 2); (1, 1, 1, 2)\rangle$. As before, for all i, i', j, j': $u_i^j \preceq u_{i'}^{j'}$ iff , $j \leq j'$ and we will use u^j to denote any u_i^j. All b^n are equal to each other because if without loss of generality $b^1 \prec b^2$ then b^1's predecessors will be stronger than b^2's predecessors and thus by strong quasi transitivity $b^2 \preceq b^1$. Again, by weak maximum transitivity, $u^{M-1} \preceq a, b \prec u^M$ and we will use b to denote any b^n. Therefore, a and b satisfy p. We can again consider a manipulation by b^1 removing an edge to a, again all b^n remain equal and as before the manipulation cannot change a and b's relative rank, and when again applying the assumption of induction we get the desired result.

By strong quasi-transitivity, profiles where all predecessors are equal are ranked $(1) \preccurlyeq (1, 1) \preccurlyeq \cdots$. When considering the result above, we conclude any two vertices should be weakly ranked according to the number of strongest predecessors they have, and by ranked IIA the tie-breaking rule must be universal.

It remains to show that vertices in $V \setminus V_s$ will be ranked equally and strictly weaker than those in V_s. Let $m \in V_s$ be a minimally ranked vertex in V_s. Consider a manipulation by m adding edges to all vertices in $V \setminus V_s$. By the above proof, all vertices in $V \setminus V_s$ will be equally ranked weaker than m. As m does not worsen its position by performing this manipulation and the internal ranking in V_s does not change we conclude that in any graph all vertices in $V \setminus V_s$ must be ranked strictly weaker than those in V_s.

We can show the vertices in $V \setminus V_s$ are ranked equally by induction on the number of edges between them. If there are no such edges, then by strong quasi-transitivity, the requirement is satisfied. Otherwise, consider an edge (v_1, v_2) such that $v_1, v_2 \in V \setminus V_s$. A manipulation by v_1 adding this edge must retain its position and thus all agents in $V \setminus V_s$ must be ranked equally.

We have shown that all vertices must be ranked according to strong count and thus the system must be a strong count system. □

5. *Relaxing the Axioms*

We shall now prove the conditions in Lemma 4.5 (and thus also in Theorem 4.3) are all necessary by showing PRSs that satisfy each three of the four conditions, but do not satisfy weak maximum transitivity. Some of these systems are quite artificial, while others are interesting and useful.

5.1. Artificial Ranking Systems.

PROPOSITION 5.1.1. *There exists a PRS that satisfies strong quasi-transitivity, ranked IIA and strong incentive compatibility, but not self-confidence nor weak maximum transitivity.*

PROOF. Let F_D^- be the PRS that ranks strictly the opposite of the distance system F_D. That is, $v_1 \preceq_{G,s}^{F_D^-} v_2 \Leftrightarrow v_2 \preceq_{G,s}^{F_D} v_1$. The proof F_D^- satisfies strong quasi transitivity, ranked IIA and strong incentive compatibility follows the proof of Proposition 3.5.1, with the following rule for ranking comparison profiles:

$$(a_1, a_2, \ldots, a_n) \preccurlyeq (b_1, b_2, \ldots, b_m) \Leftrightarrow a_1 \leq b_1.$$

F_D^- does not satisfy self-confidence, because, by definition s is weaker than all other agents, and does not satisfy weak maximum transitivity because in graph from Figure 4(a), F_D^- ranks x and y equally even though the strongest predecessor of y, which is x, is stronger than the strongest predecessor of x, which is s. □

This PRS is highly unintuitive, as the most trusted agents are the ones furthest from the source, which is by itself the least trusted.

Relaxing strong quasi-transitivity leads to a PRS that is almost trivial:

PROPOSITION 5.1.2. *There exists a PRS that satisfies self-confidence, ranked IIA and strong incentive compatibility, but not strong quasi-transitivity nor weak maximum transitivity.*

PROOF. Let F be the *PRS* that ranks for every $G = (V, E)$, for every source $s \in V$, and for every $v_1, v_2 \in V \setminus \{s\}$: $v_1 \simeq v_2 \prec s$. That is, F ranks s on the top, and all of the other agents equally. F trivially satisfies self confidence, ranked IIA and strong incentive compatibility, as s is indeed stronger than all other agents and every comparison profile is ranked equally. F does not satisfy strong quasi transitivity or weak maximum transitivity, because in a chain of vertices starting from s all except s will be ranked equally. □

5.2. RELAXING RANKED IIA. When ranked IIA is relaxed, we find a new ranking system that ranks in accordance with the distance from s, breaking ties according to the number of shortest paths from s:

Notation 5.2.1. Let $G = (V, E)$ be some directed graph and $v_1, v_2 \in V$ be some vertices, we will use $n_G(v_1, v_2)$ to denote the number of directed paths of minimum length between v_1 and v_2 in G. We will sloppily use the notations $d(v)$ and $n(v)$ to denote $d_G(s, v)$ and $n_G(s, v)$, respectively.

Definition 5.2.2. The Path Count PRS F_P is defined as follows: Given a graph $G = (V, E)$ and a source s, for all $v_1, v_2 \in V \setminus \{s\}$:

$$\begin{aligned} v_1 \preceq^{F_P}_{G,s} v_2 \Leftrightarrow\ & d_G(s, v_1) > d_G(s, v_2) \\ & \vee (d_G(s, v_1) = d_G(s, v_2) \\ & \wedge n_G(s, v_1) \leq n_G(s, v_2)) \end{aligned}$$

The Path count PRS ranks first based on distance and then based on the number of directed paths of shortest length.

Example 5.2.3. Consider again the graph in Figure 5. The vertices in the first level a and b have only one minimal length path leading to them, and so do d and e. c, however, has two, and thus $s \succ a \simeq b \succ c \succ d \simeq e$. In the next level f has the same two paths c has, but extended. g has those two paths and paths thru d and e, for a total of four paths. h has the two path thru d and e for a total of two, while i only has the one path thru e. Thus, the ranking is:

$$s \succ a \simeq b \succ c \succ d \simeq e \succ g \succ f \simeq h \succ i.$$

We shall now show that path count is indeed a result of relaxing ranked IIA:

PROPOSITION 5.2.4. *The path count PRS F_P satisfies self confidence, strong quasi transitivity and strong incentive compatibility under $\mathcal{M}_{both}$, but not ranked IIA nor weak maximum transitivity.*

PROOF. Self-confidence is trivial as $d(s) = 0 < d(v)$ for all $v \neq s$.

To prove F_P satisfies quasi-transitivity consider a graph $G = (V, E)$, a source $s \in V$ and two vertices $v_1, v_2 \in V \setminus \{s\}$. Assume for contradiction that $v_2 \prec v_1$ and there exists a 1-1 function $f : P(v_1) \mapsto P(v_2)$ such that $v \preceq f(v)$ for all $v \in P(v_1)$. By the definition of F_P: $d(v_1) \leq d(v_2)$, but

$$d(v_1) = \min_{v \in P(v_1)} d(v) + 1 \geq \min_{v \in f(P(v_1))} d(v) + 1 \geq \min_{v \in P(v_2)} d(v) + 1 = d(v_2),$$

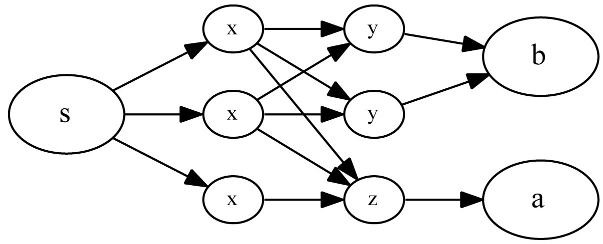

FIG. 9. Proof F_P does not satisfy ranked IIA nor weak maximum transitivity.

and thus $d(v_1) = d(v_2)$. Now,

$$\begin{aligned} n(v_1) &= \sum_{v \in P(v_1) \wedge d(v)+1=d(v_1)} n(v) \\ &\leq \sum_{v \in f(P(v_1)) \wedge d(f^{-1}(v))+1=d(v_1)} n(v) \\ &\leq \sum_{v \in P(v_2) \wedge d(v)+1=d(v_2)} n(v) = n(v_2). \end{aligned}$$

Therefore, $v_1 \preceq v_2$ in contradiction to our assumption.

For strong quasi-transitivity, assume now that $v_2 \preceq v_1$, $P(v_1) \neq \emptyset$, and there exists a 1-1 function $f : P(v_1) \mapsto P(v_2)$ such that $v \prec f(v)$ for all $v \in P(v_1)$. As above, we find that $d(v_1) = d(v_2)$. Now,

$$\begin{aligned} n(v_1) &= \sum_{v \in P(v_1) \wedge d(v)+1=d(v_1)} n(v) \\ &< \sum_{v \in f(P(v_1)) \wedge d(f^{-1}(v))+1=d(v_1)} n(v) \leq n(v_2), \end{aligned}$$

which yields $v_1 \prec v_2$ in contradiction to our assumption.

To show F_P satisfies strong incentive compatibility under $\mathcal{M}_{both}$, note that a manipulation by v cannot change $d(v)$ or $d(v')$ $\forall v' : d(v') < d(v)$. Moreover, v and its sybils cannot gain any new edges from vertices closer to v or change their internal edges. For this reason, $n(v)$ cannot increase and $n(v')$ cannot decrease for all v' s.t. $d(v') \leq d(v)$. Thus, F_P does indeed satisfy strong incentive compatibility under $\mathcal{M}_{both}$.

To show F_P does not satisfy ranked IIA nor weak maximum transitivity, consider the graph in Figure 9. F_P ranks this graph as follows: $a \prec b \prec y \prec z \prec x \prec s$. Consider the profile $\langle(2); (1, 1)\rangle$. If we compare x and y we get $(1, 1) \prec (2)$, but if we compare a and b we get $(2) \prec (1, 1)$, in violation of ranked IIA. Furthermore, the latter comparison is in violation of weak maximum transitivity, as required. □

5.3. RELAXING INCENTIVE COMPATIBILITY. When we relax incentive compatibility, we find an interesting family of PRSs that rank the agents in accordance with their in-degree, breaking ties by comparing the ranks of the strongest predecessors. We have previously presented these systems in the context of general ranking systems [Altman and Tennenholtz 2008], and we shall now adapt them to the personalized setting by explicitly letting s have the highest value.

The recursive in-degree systems work by assigning a rational number trust value for every vertex. This value 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 trust, and so on. Loops are ranked as periodic rational numbers in base $(n + 2)$ with a period the length of the loop, but only if 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 trust values can be calculated, we define these systems as follows:

Definition 5.3.1. 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 PRS with rank function* r is defined as follows: Given a graph $G = (V, E)$ and source s, the relative ranking of two vertices is based on a numeric calculation:

$$v_1 \preceq^{RID_r}_{G,s} v_2 \Leftrightarrow \text{value}_r(v_1) \leq \text{value}_r(v_2),$$

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

$$\text{value}_r(v) = \max_{\mathbf{a} \in \text{Path}(v)} \text{vp}_r(\mathbf{a}). \tag{3}$$

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

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

The path valuation function $\text{vp} : V^* \mapsto \mathbb{Q}$ defines the value to conform to a lexicographic order on in-degrees along the path, with a special exception for s:

$$\text{vp}_r(a_1, a_2, \ldots, a_m) = \frac{1}{n+2} \left[\begin{array}{c} \begin{cases} n+1 & a_1 = s \\ r(|P(a_1)|) & \text{Otherwise} \end{cases} \\ + \\ \begin{cases} 0 & m = 1 \\ \text{vp}_r(a_2, \ldots, a_m, a_2) & a_1 = a_m \wedge m > 1 \\ \text{vp}_r(a_2, \ldots, a_m) & \text{Otherwise.} \end{cases} \end{array} \right] \tag{4}$$

Note that $\text{vp}_r(a_1, a_2, \ldots, a_m)$ 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:

$$\begin{aligned}\text{vp}_r(a_1, \ldots, a_m, a_1) &= \sum_{i=0}^{\infty} \frac{1}{(n+2)^{mi}} \sum_{j=1}^{m} \frac{r(|P(a_j)|)}{(n+2)^j}\\ &= \frac{(n+2)^m}{(n+2)^m - 1} \text{vp}_r(a_1, \ldots, a_m).\end{aligned} \tag{5}$$

Further note that when the r function is constant ($r \equiv 1$), then the recursive in-degree system becomes the distance system on V_s, where the vertices in $V \setminus V_s$

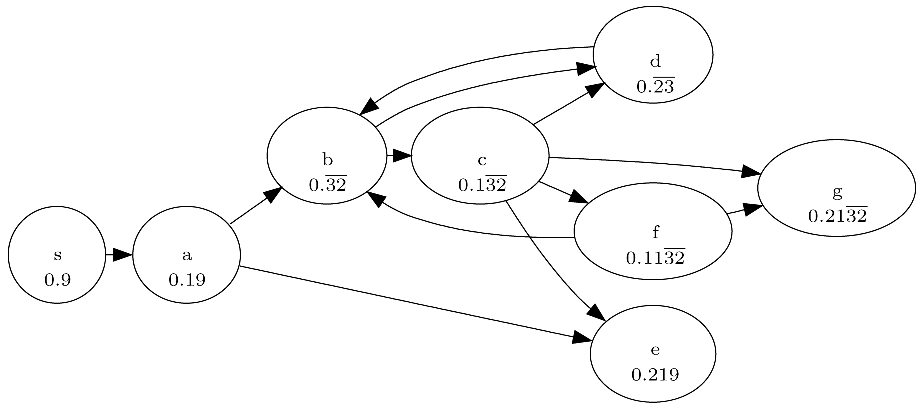

FIG. 10. Values assigned by the recursive in-degree algorithm.

are ranked weaker, and the ordering among them is set according to the length of the longest path (simple or not) leading to the vertex.

Example 5.3.2. An example of the values assigned for a particular graph when r is the identity function is given in Figure 10. As $n = 8$, and the definition in (4) is based on recursive division by $n + 2$, the trust values are decimals which consist of a concatenation of in-degrees along the maximal path, or 9 in the case of s.

The value 0.9 is assigned to s by the first case in (4). The value for a arises from the path (s, a) and the last case in (4), where the recursive call gives the value of s (0.9). This is added to $r(|P(a)|) = 1$ and divided by 10, giving the result 0.19.

The values of b and d arise from a loop consisting of these vertices. Applying the middle case in (4), we have the equations

$$\begin{aligned} \text{value}_r(b) &= \text{vp}_r(b, d, b) = \frac{1}{10}[3 + \text{vp}_r(d, b, d)] \\ \text{value}_r(d) &= \text{vp}_r(d, b, d) = \frac{1}{10}[3 + \text{vp}_r(b, d, b)] \end{aligned}$$

By using (5), we get the periodic decimals seen in Figure 10. The rest of the values are similarly obtained by the last case in (4).

Note that even though there are several loops in the graph, the other loops are not on a maximal path as defined above. An algorithm for efficiently computing recursive-indegree is given in Altman and Tennenholtz [2008].

We shall now show that indeed all the axioms from Theorem 4.3 are satisfied except incentive compatibility.

PROPOSITION 5.3.3. *Let $r : \mathbb{N} \mapsto \mathbb{N}$ be a 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 self-confidence, strong quasi-transitivity and ranked IIA. If r is not constant,*[5] *then the recursive in-degree system further does not satisfy*

[5] If r is constant, the system still does not satisfy strong incentive compatibility under either $\mathcal{M}_{out}$ or $\mathcal{M}_{sybil}$, but only if we allow vertices that have no path from s.

weak maximum transitivity nor strong incentive compatibility under either $\mathcal{M}_{out}$ or $\mathcal{M}_{sybil}$.

PROOF. We will prove that in the entire graph (not just V_s) 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:

$$\mathbf{a} \preccurlyeq \mathbf{b} \;\Leftrightarrow\; (k = 0) \vee (r(k) < r(l)) \vee [(r(k) = r(l)) \wedge (a_k \leq b_l)] \,.$$

Note that this ranking of comparison profiles also implies strong quasi-transitivity. To show comparison profiles are ranked as such, we will prove that

$$\text{value}_r(v) \;=\; \begin{cases} 0 & v \neq s \wedge P(v) = \emptyset \\ \frac{n+1}{n+2} & v = s \\ \frac{1}{n+2}\left[r(|P(v)|) + \max_{p \in P(v)} \text{value}_r(p)\right] & \text{Otherwise} \end{cases} \tag{6}$$

and note that $0 \leq \text{value}_r(v) \leq \frac{n+1}{n+2}$, and thus vertices other than s are ordered first by $r(|P(v)|)$ and then by $\max_{p \in P(v)} \text{value}_r(p)$, as required. Moreover, self-confidence is satisfied because for all $v \neq s$: $\text{value}_r(v) < \frac{n+1}{n+2}$.

The two edge cases are trivial, we shall now concentrate on the primary case in (6). Let $v \in V \setminus \{s\}$ be some vertex where $P(v) \neq \emptyset$. Denote $\text{Path}'(p, v)$ as the set of almost-simple directed paths to p stopping at s, which do not pass through v unless immediately looping back to p:

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

Now we see that:

$$\begin{aligned} \text{value}_r(v) \;&=\; \max_{\mathbf{a} \in \text{Path}(v)} \text{vp}_r(\mathbf{a}) = \\ &=\; \frac{1}{n+2}\left[\begin{array}{l} r(|P(v)|) + \max_{(v=a_1,\ldots,a_m) \in \text{Path}(v)} \\ \begin{cases} \text{vp}_r(a_2, \ldots, a_m, a_2) & a_1 = a_m \wedge m > 1 \\ \text{vp}_r(a_2, \ldots, a_m) & \text{Otherwise.} \end{cases} \end{array}\right] && (7) \\ &=\; \frac{1}{n+2}\Big[r(|P(v)|) + \max_{p \in P(v)} \max_{\mathbf{a} \in \text{Path}'(p,v)} \text{vp}_r(\mathbf{a})\Big] && (8) \\ &=\; \frac{1}{n+2}\Big[r(|P(v)|) + \max_{p \in P(v)} \max_{\mathbf{a} \in \text{Path}(p)} \text{vp}_r(\mathbf{a})\Big] \\ &=\; \frac{1}{n+2}\Big[r(|P(v)|) + \max_{p \in P(v)} \text{value}_r(p)\Big]. \end{aligned}$$

To show that the equality (8) holds, assume for contradiction that there exists $p \in P(v)$ and $\mathbf{a} \in \text{Path}(p)$ such that

$$\text{vp}_r(\mathbf{a}) > \max_{p' \in P(v)} \max_{\mathbf{a}' \in \text{Path}'(p',v)} \text{vp}_r(\mathbf{a}'). \tag{9}$$

From $\mathbf{a} \in \text{Path}(p) \setminus \text{Path}'(p, v)$, 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 $j < i$

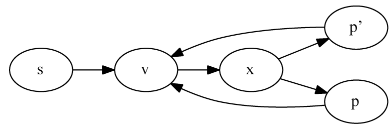

$$\mathbf{a} = (p, x, v, p', x)$$
$$\mathbf{b} = (p, x, v, p)$$
$$\mathbf{c} = (p', x, v, p')$$

FIG. 11. Example of paths from the proof of Proposition 5.3.3.

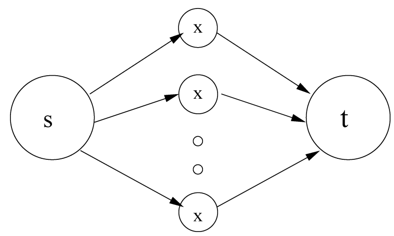

FIG. 12. Graph from proof that Recursive In-degree is not incentive compatible.

or $(p' = a_{i+1}, \ldots, a_m)$ otherwise. An example of such paths is given in Figure 11. Note that $\mathbf{b} \in \text{Path}'(p, v)$ and $\mathbf{c} \in \text{Path}'(p', v)$, where $p, p' \in P(v)$. Now, note that

$$\text{vp}_r(\mathbf{a}) = \frac{(n+2)^j - 1}{(n+2)^j} \text{vp}_r(\mathbf{b}) + \frac{1}{(n+2)^j} \text{vp}_r(\mathbf{c}),$$

and thus $\text{vp}_r(\mathbf{a})$ must be between $\text{vp}_r(\mathbf{b})$ and $\text{vp}_r(\mathbf{c})$, in contradiction to assumption (9).

We shall now prove that recursive in-degree is not incentive compatible under $\mathcal{M}_{out}$ or $\mathcal{M}_{sybil}$ and does not satisfy weak maximum transitivity. Let $i \in \mathbb{N}$ be the minimum number such that $r(i) > 1$. Consider the graph G in Figure 12, where there are i vertices labeled x. This graph is ranked $x \prec t \prec s$, where x refers to all vertices labeled x. Weak maximum transitivity is not satisfied because $x \prec t$ even though $s \succ x$. Let x' be one of the vertices labeled x. It can perform a manipulation in $\mathcal{M}_{out}$ by removing its edge to t, and thus changing the ranking to $x \simeq x' \simeq t \prec s$. It can also perform a manipulation in $\mathcal{M}_{sybil}$ by creating i additional sybils of itself and creating a complete clique thus changing the ranking to $x \prec v \simeq x' \simeq t \prec s$, where v are the new vertices involved in the manipulation. □

For an extensive study of the recursive in-degree system in the context of general ranking systems see Altman and Tennenholtz [2008].

6. *Concluding Remarks*

We have presented a method for the evaluation of personalized ranking systems by using axioms adapted from the ranking systems literature, and have evaluated existing and new personalized ranking systems according to these axioms. As most existing PRSs do not satisfy these axioms, we have presented several new and practical personalized ranking systems that satisfy subsets, or indeed all, of these axioms. We argue that these new ranking systems have a more solid theoretical basis, and thus may very well be successful in practice.

Furthermore, we have proven a representation theorem for the Strong Count ranking systems, which are the only systems that satisfy all our axioms.

This study is far from exhaustive. Further research is due in formulating new axioms, and proving representation theorems for the various PRSs suggested in this paper. An additional avenue for research is modifying the setting in order to accommodate for more elaborate input such as trust/distrust relations or numerical trust ratings, as seen in some existing personalized ranking systems used in practice.

REFERENCES

ALTMAN, A., AND TENNENHOLTZ, M. 2005. Ranking systems: the PageRank axioms. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC'05)*. ACM, New York, 1–8.

ALTMAN, A., AND TENNENHOLTZ, M. 2006. Quantifying incentive compatibility of ranking systems. In *Proceedings of AAAI-06*. AAAI, Menlo Park, CA.

ALTMAN, A., AND TENNENHOLTZ, M. 2007. Incentive compatible ranking systems. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-07)*. ACM, New York, 84.

ALTMAN, A., AND TENNENHOLTZ, M. 2008. Axiomatic foundations for ranking systems. *J. Artif. Intell. Res.*, *31*, 1(Jan.), 473–495.

ANDERSEN, R., BORGS, C., CHAYES, J., FEIGE, U., FLAXMAN, A., KALAI, A., MIRROKNI, V., AND TENNENHOLTZ, M. 2008. Trust-based recommendation systems: an axiomatic approach. In *Proceeding of the 17th International Conference on World Wide Web (WWW'08)*. ACM, New York, 199–208.

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

AVESANI, P., MASSA, P., AND TIELLA, R. 2005. A trust-enhanced recommender system application: Moleskiing. In *Proceedings of the ACM Symposium on Applied Computing (SAC'05)*. ACM, New York, 1589–1593.

BOGOMOLNAIA, A., MOULIN, H., AND STONG, R. 2005. Collective choice under dichotomous preferences. *J. Econ. Theory 122*, 2 (June), 165–184 (available at http://ideas.repec.org/a/eee/jetheo/v122y2005i2p165-184.html).

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.

CALLAS, J., DONNERHACKE, L., FINNEY, H., AND THAYER, R. 1998. RFC 2440: OpenPGP message format. Status: PROPOSED STANDARD.

CHENG, A., AND FRIEDMAN, E. 2005. Sybilproof reputation mechanisms. In *Proceeding of the ACM SIGCOMM Workshop on Economics of Peer-to-Peer Systems (P2PECON'05)*. ACM, New York, 128–132.

CHIEN, S., DWORK, C., KUMAR, R., SIMON, D. R., AND SIVAKUMAR, D. 2003. Link evolutions: Analysis and algorithms. *Internet. Math 1*, 377–304.

DASH, R. K., RAMCHURN, S. D., AND JENNINGS, N. R. 2004. Trust-based mechanism design. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'04)*. IEEE Computer Society Press, Los Alamitos, CA, 748–755.

DELLAROCAS, C. 2003. Efficiency through feedback-contingent fees and rewards in auction marketplaces with adverse selection and moral hazard. In *Proceedings of the 3rd ACM Conference on Electronic Commerce (EC-03)*. ACM, New York, 11–18.

DUTTA, B., JACKSON, M. O., AND LE BRETON, M. 2001. Strategic candidacy and voting procedures. *Econometrica 69*, 4, 1013–37.

FISHBURN, P. C. 1978. Axioms for approval voting: Direct proof. *J. Econ. Theory 19*, 1 (Oct.), 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.

GUHA, R., KUMAR, R., RAGHAVAN, P., AND TOMKINS, A. 2004. Propagation of trust and distrust. In *Proceedings of the 13th International Conference on World Wide Web (WWW'04)*. ACM, New York, 403–412.

GYÖNGYI, Z., GARCIA-MOLINA, H., AND PEDERSEN, J. 2004. Combating web spam with trustrank. In *Proceedings of the 13th International Conference on Very Large Data Bases (vldb'04)*. VLDB Endowment, 576–587.

HAVELIWALA, T., KAMVAR, S., AND JEH, G. 2003. An analytical comparison of approaches to personalizing PageRank. Tech. rep., Stanford University, Stanford, CA.

JEH, G., AND WIDOM, J. 2003. Scaling personalized web search. In *Proceedings of the 12th International Conference on World Wide Web (WWW'03)*. ACM, New York, 271–279.

KLEINBERG, J. M. 1999. Authoritative sources in a hyperlinked environment. *JACM 46*, 5, 604–632.

LEE, H. C., AND BORODIN, A. 2003. Perturbation of the hyper-linked environment. In *Proceedings of the 9th Annual International Conference on Computing and Combinatorics (COCOON 2003)*. Lecture Notes in Computer Science, vol. 2697. Springer-Verlag, Berlin, Germany, 272–283.

LEMPEL, R., AND MORAN, S. 2005. Rank-stability and rank-similarity of link-based web ranking algorithms in authority-connected graphs. *Inf. Retr. 8*, 2, 245–264.

LEVIEN, R. 2002. Attack Resistant Trust Metrics. Ph.D. dissertation. University of California, Berkeley, CA.

MASSA, P., AND AVESANI, P. 2005. Controversial users demand local trust metrics: An experimental study on epinions.com community. In *Proceedings of AAAI-05*. AAAI, Menlo Park, CA, 121–126.

NG, A. Y., ZHENG, A. X., AND JORDAN, M. I. 2001. Stable algorithms for link analysis. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'01)*. ACM, New York, 258–266.

PAGE, L., BRIN, S., MOTWANI, R., AND WINOGRAD, T. 1998. The PageRank citation ranking: Bringing order to the web. Tech. rep., Stanford Univ. Stanford, CA.

PALACIOS-HUERTA, I., AND VOLIJ, O. 2004. The measurement of intellectual influence. *Econometrica 73*, 3.

PENNOCK, D. M., HORVITZ, E., AND GILES, C. L. 2000. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of the 17th National Conference on Artificial Intelligence and 12th Conference on Innovative Applications of Artificial Intelligence*. AAAI Press/The MIT Press, Cambridge, MA, 729–734.

RESNICK, P., KUWABARA, K., ZECKHAUSER, R., AND FRIEDMAN, E. 2000. Reputation systems. *Commun. ACM 43*, 12, 45–48.

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.

SATTERTHWAITE, M. 1975. Strategy proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *J. Econ. Theory 10*, 187–217.

TENNENHOLTZ, M. 2004. Reputation systems: An axiomatic approach. In *Proceedings of the 20th Conference on Uncertainity in Artificial Intelligence (UAI-04)*. AVAI Press, 544–551.

WU, B., AND DAVISON, B. 2005. Identifying link farm spam pages. In *Proceedings of the 14th International World Wide Web Conference, Industrial Track*. ACM, New York.

WU, B., GOEL, V., AND DAVISON, B. D. 2006. Propagating trust and distrust to demote web spam. In *Proceedings of the WWW'06 Workshop on Models of Trust for the Web (MTW'06)*. ACM, New York.

RECEIVED MARCH 2007; REVISED FEBRUARY 2008 AND JANUARY 2009; ACCEPTED JANUARY 2010