

Trust-Based Recommendation Systems: an Axiomatic Approach

Reid Andersen*
Abraham Flaxman*

Christian Borgs*
Adam Kalai^{§ †}

Jennifer Chayes*
Vahab Mirrokni*

Uriel Feige^{† ‡}
Moshe Tennenholtz^{† ‡}

ABSTRACT

High-quality, personalized recommendations are a key feature in many online systems. Since these systems often have explicit knowledge of social network structures, the recommendations may incorporate this information. This paper focuses on networks that represent trust and recommendation systems that incorporate these trust relationships. The goal of a trust-based recommendation system is to generate personalized recommendations by aggregating the opinions of other users in the trust network.

In analogy to prior work on voting and ranking systems, we use the axiomatic approach from the theory of social choice. We develop a set of five natural axioms that a trust-based recommendation system might be expected to satisfy. Then, we show that no system can simultaneously satisfy all the axioms. However, for any subset of four of the five axioms we exhibit a recommendation system that satisfies those axioms. Next we consider various ways of weakening the axioms, one of which leads to a unique recommendation system based on random walks. We consider other recommendation systems, including systems based on personalized PageRank, majority of majorities, and minimum cuts, and search for alternative axiomatizations that uniquely characterize these systems.

Finally, we determine which of these systems are incentive compatible, meaning that groups of agents interested in manipulating recommendations can not induce others to share their opinion by lying about their votes or modifying their trust links. This is an important property for systems deployed in a monetized environment.

Categories and Subject Descriptors

F.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

*Microsoft Research, Redmond, WA.

[†]Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, Israel.

[‡]Part of this research conducted while visiting MSR.

[§]School of Computer Science, Georgia Institute of Technology, Atlanta, GA. Partly supported by NSF SES-0734780.

[¶]Faculty of Industrial Engineering and Management, Technion–Israel Institute of Technology, Haifa, Israel.

General Terms

Algorithms, Theory, Economics

Keywords

Recommendation systems, Reputation systems, Axiomatic approach, Trust networks

1. INTRODUCTION

Ranking, reputation, recommendation, and trust systems have become essential ingredients of web-based multi-agent systems (e.g. [16, 22, 8, 25, 11]). These systems aggregate agents' reviews of products, services, and each other, into valuable information. Notable commercial examples include Amazon and E-Bay's recommendation and reputation systems (e.g. [21]), Google's page ranking system [19], and the Epinions web of trust/reputation system (e.g. [18]).

This paper is concerned particularly with the setting where there is a single item of interest (e.g., a product, service, or political candidate). A subset of the agents have prior opinions about this item. Any of the remaining agents might desire to estimate whether or not they would like the item, based on the opinions of others. In the offline world, a person might first consult her friends for their recommendations. In turn, the friends, if they do not have opinions of their own, may consult their friends, and so on. Based on the cumulative feedback the initial consulter receives, she might form her own subjective opinion. An automated trust-based recommendation system aims to provide a similar process to produce high-quality *personalized* recommendations for agents.

We model this setting as an annotated directed graph in which some of the nodes are labeled by votes of + and −. Here a node represents an agent, and an edge directed from a to b represents the fact that agent a trusts agent b . A subset of the nodes are labeled by + or − votes, indicating that these nodes have already formed opinions about the item under question. Based on this input, a recommendation system must output a recommendation for each unlabeled node. We call such an abstraction a *voting network* because it models a variety of two-candidate voting systems, where the candidates are + and −. For an example, consider a directed star graph where a single root node points to n agents with labels, which models a committee making a recommendation to the root node. In that setting, majority and consensus are two common voting rules. For another example, The U.S. presidential voting system can be modeled as a more complicated digraph, where the root points

to nodes representing the members of the electoral college, and the electoral college nodes point to nodes representing the voters in the state or congressional district that they represent.

A multitude of recommendation systems have been proposed and implemented, and many fit in to the network-based framework described above. This raises the question of how to determine the relative merits of alternative approaches to providing trust-based recommendations. The task of comparing recommendation systems is complicated by the difficulty of producing an objective measure of recommendation quality.

We begin with an impossibility theorem: for a small, natural set of axioms, there is no recommendation system simultaneously consistent with all axioms in the set. However, for any proper subset of the axioms there exists a recommendation system that satisfies all axioms in the subset. We consider two ways past this negative result, both by replacing the *transitivity* axiom (defined below). We prove that there are recommendation systems consistent with both new sets of axioms. We also show that when one of these new sets is augmented with an additional axiom, the resulting set of axioms is uniquely satisfied by a recommendation system based on random walks.

We also consider the descriptive approach, in which we characterize existing (acyclic) systems, like simple committees and the U.S. presidential elections, by a simple majority axiom. We generalize this to an axiom that leads to a unique “minimum cut” system on general undirected (possibly cyclic) graphs.

Prior work on personalized ranking systems has produced a ranking system called *personalized PageRank* [13]. This system can be translated into a recommendation system by augmenting it to handle votes, creating a natural hybrid of a ranking and recommendation system. The details are discussed in Section 3.4.

We define a notion of incentive compatibility for recommendation systems. This is important when designing systems for deployment in monetized settings, because, as experience has shown, self-interested agents will not respect the rules of the system when there is money to be made by doing otherwise. We find that all of the recommendation systems for which we provide a characterizing set of axioms turn out to be incentive compatible, including the random walk system, majority of majorities system, and minimum cut system. In contrast, the personalized PageRank system and various other natural systems are not incentive compatible.

For simplicity, this paper focuses on the case of unweighted multigraphs and binary votes. Most of our observations carry over to the more general cases of weighted graphs and fractional votes.

The rest of this paper is organized as follows. The next section contains related work. Section 1.2 is a brief summary of our axioms, algorithms, and results. Some formal definitions and notations appear in Section 2, and in Section 3 we present some recommendation systems. In Section 4 we formally define our axioms. Section 5 provides the outline of the proofs of our results. Section 6 shows that our systems are incentive compatible.

1.1 Related Work

There are several ways to study recommendation systems.

Standard evaluation tools include simulations and field experiments (e.g. [7, 21, 15]). In addition, one may also consider computational properties of suggested systems.

Our work builds on previous work on axiomatizations of ranking systems. The literature on the axiomatic approach to ranking systems deals with both global ranking systems [1, 2, 24, 9, 25, 4, 5] and personalized ranking systems [7, 10, 3, 17]. Personalized ranking systems are very close to trust-based recommendation systems. In such systems, agents rank some of the other agents. Then an aggregated ranking of agents, personalized to the perspective of a particular agent, is generated based on that information. However, previous studies on the axiomatic approach have not been concerned with situations where the participants share reviews or opinions on items of interest which are not the other agents in the system.

Many existing recommendation systems are based on collaborative filtering (CF), which is a completely different approach than the trust-based systems considered in this paper. An axiomatic approach to analyzing CF systems appears in [20]. Combining trust-based and CF approaches is a direction of current research [23].

1.2 Overview of Results

We model a *voting network* by a partially labeled directed multigraph whose nodes represent agents. A subset of the nodes, called *voters*, is labeled with votes of + and -. The remaining nodes are *nonvoters*. The recommendation system must assign, to each *source* nonvoter, a *recommendation* in $\{-, 0, +\}$.¹

Below we informally summarize our axioms. Many of them are illustrated in Figure 1. We caution that our aim is only to succinctly convey the spirit of the axioms; formal definitions are found in Section 2.

1. **Symmetry.** Isomorphic graphs result in corresponding isomorphic recommendations (anonymity), and the system is also symmetric with respect to + and - votes (neutrality).
2. **Positive response.** If a node’s recommendation is 0 and an edge is added to a + voter, then the former’s recommendation becomes +.
3. **Independence of Irrelevant Stuff (IIS).** A node’s recommendation is independent of agents not reachable from that node. Recommendations are also independent of edges leaving *voters*.
4. **Neighborhood consensus.** If a nonvoter’s neighbors *unanimously* vote +, then the recommendation of *other* nodes will remain unchanged if that node instead becomes a + voter.

If, in a particular graph, a source node is recommended +, then we say that the source trusts the set of agents that voted + more than those that voted -. The following axiom states that this relation should be transitive, as we imagine varying the votes of various subsets of agents.

5. **Transitivity.** For any graph (N, E) and disjoint sets $A, B, C \subseteq N$, for any source s , if s trusts A more than B , and s trusts B more than C , then s trusts A more than C .

¹It is easy to see, e.g., a three-node network, that symmetry (Axiom 1) is impossible with \pm recommendations alone.




Figure 1: Illustrative voting networks. Labels \pm indicate votes, not recommendations. a) IIS: Node s 's recommendation does not depend on any of the dashed node or dashed edges, since we ignore unreachable nodes as well as edges out of voters. b) Neighborhood consensus: Node v can be assigned $+$ vote. c) If the recommendation for s is $+$ then we say s trusts $\{v\}$ more than $\{u\}$. d) Trust propagation: The dashed edges the upper part can be removed and replaced by dashed edges in the lower part. e) Scale invariance: Edges leaving s are tripled without consequence. f) No groupthink: The three nonvoting nodes cannot all be given $+$ recommendations.

THEOREM 1. *Axioms 1-5 are inconsistent. Any proper subset of these axioms is satisfied by some recommendation system.*

Instead of transitivity, we consider the following two axioms. Similar axioms have been used in the axiomatization of PageRank in the context of ranking systems [2].

6. **Trust propagation.** Suppose node u trusts nonvoter v . Suppose that u has k edges to v and v has k edges to other nodes. Then the edges from u to v can be replaced directly by edges from u to the nodes that v trusts without changing any resulting recommendations.
7. **Scale invariance.** Duplicating the outgoing edges of a node does not change recommendations.

We specify a *random walk algorithm*, a (deterministic) algorithm that computes a recommendation for node s by considering a (hypothetical) random walk in the directed graph that starts at node s and follows outgoing edges, uniformly at random, until the first voter is reached. Its recommendation for s is based on whether it is *more likely* that the first voter reached votes $+$ or $-$. The system computed by the random walk algorithm is called the *Random Walk recommendation system* (RW).

THEOREM 2. *Axioms 1-4 and 6-7 are satisfied uniquely by RW.*

As we will see in the proof of Theorem 1, the transitivity axiom and the IIS axiom are hard to reconcile because IIS implies the edges out of the voting nodes do not matter while transitivity implies that the sets in the trust graph must obey a certain relation regardless of who is voting. A weaker version of transitivity which does not conflict with IIS is the following:

- 5'. **Sink Transitivity.** For any graph (N, E) and any disjoint sets $A, B, C \subseteq N$ for which A, B, C contain only vertices with out-degree 0, for any source s , if s trusts A more than B and s trusts B more than C , then s trusts A more than C .

THEOREM 3. *Axioms 1-4 and 5' are satisfied by RW.*

We also consider axioms which lead uniquely to known recommendation systems:

8. **Majority.** The recommendation of a node should be equal to the majority of the votes and recommendations of its trusted neighbors.
9. **No groupthink.** Suppose a set of nonvoters unanimously have the same nonzero recommendation. Then their recommendation should equal the majority of their trusted (external) neighbors' votes and recommendations.

The majority axiom represents a reasonable semantics — an agent might like to wait for its trusted neighbors to receive a recommendation and then take a simple majority. However, this axiom alone still permits a large clique of nonvoters to all have positive recommendations when they only point to external agents with negative recommendations (see Figure 1f).

The no-groupthink axiom is a natural extension to larger sets. It *implies* the majority axiom when one considers just singleton sets.

Unfortunately, on general directed graphs axiom 9 is inconsistent. However, it is a statement about a single graph G , so we can consider it on limited classes of graphs. Two interesting classes of graphs are directed acyclic graphs and undirected graphs, where axioms 8 and 9 lead uniquely to two interesting solutions. These are the majority-of-majorities and minimum-cut systems are defined in Section 3.

THEOREM 4. (a) *Axiom 8 on a rooted DAG implies the majority-of-majorities system.* (b) *Axiom 9 on an undirected graph implies the min-cut recommendation system.*

2. NOTATION AND DEFINITIONS

Following the motivation provided in the previous section, we now formally define the basic setting of a *trust-based recommendation system*. In the remainder of the paper, we refer to such systems simply as *recommendation systems*, for brevity.

DEFINITION 1. A **voting network** is a directed annotated multigraph $G = (N, V_+, V_-, E)$ where N is a set of nodes, $V_+, V_- \subseteq N$ are disjoint subsets of positive and negative voters, and $E \subseteq N^2$ is a multiset of edges with parallel edges allowed but no self-loops.

When V_+ and V_- are clear from context, we denote the set of **voters** by $V = V_+ \cup V_-$ and the set of **nonvoters** by $\bar{V} = N \setminus V$.

DEFINITION 2. A **recommendation system** R takes as input a voting network G and source $s \in \bar{V}$ and outputs recommendation $R(G, s) \in \{-, 0, +\}$.

For convenience, we will use $R_+(G), R_-(G)$, and $R_0(G)$ to denote the set of sources to which R gives a particular recommendation, i.e. $R_+ = \{s \in \bar{V} \mid R(G, s) = +\}$. Also, we define $R(G) = \langle R_+(G), R_-(G), R_0(G) \rangle$.

We denote by $\text{sgn} : \mathbb{R} \rightarrow \{-, 0, +\}$ the function that computes the sign of its input. We denote by $\text{Pred}_E(v)$ and $\text{Succ}_E(v)$ the multisets of nodes that point to v and that v points to, respectively (where, for example, u appears in $\text{Pred}_E(v)$ with multiplicity equal to the number of arcs (u, v) in multiset E).

Given a multiset of recommendations, $S \subseteq \{-, 0, +\}$, we define the *majority* $\text{MAJ}(S)$ to be: $+$ if more than half the elements of S are $+$; $-$ if more than half of S are $-$; and 0 otherwise.

3. SYSTEMS AND ALGORITHMS

3.1 Random Walk System (RW)

We first describe a recommendation system based on random walks. Given a voting network $G = (N, V_+, V_-, E)$ and a source $s \in \bar{V}$, the random walk system simulates the following: start a walker at node s and, at each step, choose an outgoing edge uniformly at random and follow it to the destination node. Continue this random walk until a node with a \pm vote is reached, or until a node with no outgoing edges is reached (note this walk may reach a node whose strongly connected component contains no voters, and thus never terminate). Let p_s be the probability that the random walk terminates at a node with positive vote and q_s be the probability that the random walk terminates at node with negative vote. Let $r_s = p_s - q_s$. The **random walk recommendation system** recommends $\text{sgn}(r_s)$ to s .

The algorithm given in Figure 2 correctly computes recommendations as defined by the RW system. To see this, first note that, for any node that cannot reach a voter, the recommendation must be 0 in both the RW system and the RW algorithm. Next, probabilistic arguments show that the equations defined must be obeyed by the random walk. Finally, the uniqueness of the solution follows from the fact that, with probability 1 , the random walk will reach a vertex in $S \cup V_+ \cup V_-$. Also note that the algorithm can be implemented efficiently and in polynomial time.

As stated in Theorem 2 (Section 1.2), the RW recommendation system is the unique system which satisfies axioms 1-4 and 6-7. The formal definitions of these axioms appear in Section 4 and the proof of Theorem 2 is sketched in Section 5 and provided in full in Appendix B. Theorem 3 states that RW satisfies 1-4 and $5'$. The proof of this is omitted.

Input: $G = (N, V_+, V_-, E)$, $s \in \bar{V}$.

Output: recommendation $\in \{-, 0, +\}$.

1. Let $S \subseteq \bar{V}$ be the set of nonvoters that cannot reach any voter.
2. For each $v \in N$, create a variable $r_v \in \mathbb{R}$. Solve the following from r_v :

$$r_v = \begin{cases} 0, & \text{if } v \in S \\ 1, & \text{if } v \in V_+ \\ -1, & \text{if } v \in V_- \\ \frac{\sum_{w \in \text{Succ}_E(v)} r_w}{|\text{Succ}_E(v)|}, & \text{otherwise} \end{cases}$$

3. Output $\text{sgn}(r_s)$.
-

Figure 2: The random walk algorithm. (Recall that $V = V_+ \cup V_-$ is the set of voters and $\bar{V} = N \setminus V$ is the set of nonvoters.)

3.2 Majority-of-Majorities (MoM)

The system in this section and in the next seem quite different at first glance, but both derive from the same single axiom.

In this section, we present a system that is well defined only when the graph underlying the voting network is a Directed Acyclic Graph (DAG). Nodes in a finite DAG can always be partitioned into a finite number of *levels*. In level 0 , nodes have out-degree 0 . In each level $i \geq 1$, nodes have edges only to nodes in level $j < i$.

The **Majority-of-majorities** system assigns a recommendation as follows: each nonvoter that is a sink (i.e. in level 0) receives recommendation 0 ; each voter in level i receives a recommendation equal to the majority of the *recommendations and votes* of its outgoing neighbors (where multiple edges are counted multiple times). This can be computed recursively by an efficient algorithm. Recall that we use the definition of Majority from Section 2, which is conservative in the sense that in order to have a non-zero recommendation, there must be a strict majority of matching recommendations.

As stated in Theorem 4a, the majority-of-majorities recommendation system is the unique system which satisfies axiom 8 when the voting network is a DAG. The proof of Theorem 4 is sketched in Section 5 and provided in full in Appendix C.

3.3 Minimum Cut System (min-cut)

Let $G = (N, V_+, V_-, E)$ be a voting network. Let $E' \subseteq E$ be the set of edges in E that originate at nonvoters, i.e., eliminate edges out of voters. We say that $C \subseteq E'$ is a (V_+, V_-) -cut of G if there is no path from V_+ to V_- using edges in $E' \setminus C$. We say that C is a min-cut of G if its *size* $|C|$ is minimal among all (V_+, V_-) -cuts.

The **minimum cut system** can be defined as follows. The recommendation of a source s is $+$ if in *all* min-cuts there is a path from s to V_+ among edges in $E' \setminus C$. The recommendation is $-$ if all min-cuts leave a path from s to V_- . Otherwise, the recommendation is 0 .

This may be computed as follows: first, compute a min-cut C in G . Next, consider network G_+ which is formed

from G by adding an edge from source s to a $+$ voter, and compute C_+ in G_+ . If $|C| < |C_+|$ then the recommendation is $-$. Otherwise, consider network G_- , formed from G by adding an edge from s to a $-$ voter. For a min-cut C_- of G_- , if $|C| < |C_-|$ then the recommendation for s is $+$. Otherwise, the recommendation for s is 0 . This can be computed in polynomial time by repeatedly using a polynomial-time algorithm for (s, t) -minimum cut. To see that it correctly computes the min-cut recommendation, note that if s is connected to V_+ in all min-cuts then adding an edge from s to a $-$ voter will create a path from V_- to V_+ in all min-cuts. This will increase the min-cut cost by 1. On the other hand, if there is a cut of G where s is not connected to V_+ , then this edge set will still be a cut in G_- so the min-cut cost will not increase. Similarly, $|C| < |C_+|$ if and only if s is connected to V_- in all min-cuts of G . In the remaining case, the sizes of all three min-cuts will be the same because there are some min-cuts in which s is connected to V_- and some in which s is connected to V_+ .

As stated in Theorem 4b, the min-cut recommendation system is the unique system which satisfies axiom 8 on undirected graphs. The proof of Theorem 4 is sketched in Section 5 and provided in full in Appendix C.

3.4 Personalized PageRank System

In designing a recommendation system, we can consider systems based on aggregating the *level of trust* a node has for other nodes. The idea is to define a trust level w_{uv} for any two nodes u and v in the network, and to compute the recommendation for a nonvoter u by comparing two values: $W_+(u) = \sum_{v \in V_+} w_{uv}$ and $W_-(u) = \sum_{v \in V_-} w_{uv}$.

A natural way to capture the level of trust in a network is to apply a personalized ranking system such as personalized PageRank (as introduced by Haveliwala in [13]). The personalized PageRank of node v for a node u is denoted by $\text{ppr}_\alpha(u \rightarrow v)$, and is defined to be the probability mass at node v in the stationary distribution of a biased random walk $\mathcal{R}_\alpha(u)$ with restart probability $\alpha \in (0, 1)$. The random walk $\mathcal{R}_\alpha(u)$ is defined as follows: at each step, with probability α , we restart the random walk at node u ; and with probability $1 - \alpha$, we go uniformly to one of the outgoing neighbors, i.e., from a node w we go to one of the neighbors of w with probability $\frac{1-\alpha}{\text{out-degree}(w)}$; if there are no outgoing neighbors, we restart the walk at node u . Note that, unlike the RW system in Section 3.1, this random walk continues regardless of whether the vertex it is currently visiting is a voter or nonvoter. The personalized PageRank values can be computed efficiently by simulating this random walk. For details of computation and extensions, see [13, 14].

Using this notation, the personalized PageRank recommendation system is as follows: Given a voting network $G = (N, V_+, V_-, E)$, and a parameter α , we compute the personalized PageRank $\text{ppr}_\alpha(u \rightarrow v)$ for all pairs of nodes. For a nonvoter s , we compute $W_+(s) = \sum_{v \in V_+} \text{ppr}_\alpha(s \rightarrow v)$, and $W_-(s) = \sum_{v \in V_-} \text{ppr}_\alpha(s \rightarrow v)$, and we set $R(G, s) = \text{sgn}(W_+ - W_-)$.

It is not hard to see that the personalized PageRank system satisfies the axioms symmetry, positive response, and transitivity, but it does not satisfy the axioms IIS, and neighborhood consensus. The proofs of these statements are omitted due to space constraints.

4. AXIOMS

We are now ready to consider various properties of recommendation systems as candidate axioms. These properties are motivated by related literature on social choice and ranking systems, and also by the machinery used in practical recommendation systems. Similar to other axiomatic studies, the choice of axioms is to some extent arbitrary, and other sets of axioms are possible. Nevertheless, we believe that any one of our axioms by itself does capture a desirable property for recommendation systems, and that the study of the combination of these axioms leads to informative insights and interesting algorithms.

The first property, symmetry, is purely structural. The symmetry axiom contains two parts, anonymity and neutrality. Anonymity means that the names of the agents do not matter for the source node; all that matters is the structure of the trust graph and the votes provided. Neutrality means that the values $+/-$ are arbitrary.

AXIOM 1. (Symmetry) *Let $G = (N, V_+, V_-, E)$ be a voting network. Anonymity: For any permutation $\pi : N \rightarrow N$, let G' , be the isomorphic voting network under π . Then $R_+(G') = \pi(R_+(G))$ and $R_-(G') = \pi(R_-(G))$. Neutrality: Also, let $G'' = (N, V_-, V_+, E)$ be the voting network where the sets of voters V_- and V_+ have been interchanged. Then $R_+(G) = R_-(G'')$ and $R_-(G) = R_+(G'')$.*

The next axiom is a classic social choice axiom. It states that if a node s has recommendation 0 (or $+$) and an additional $+$ -voter is added to the network along with an edge from s to the new node, then s 's new recommendation should be $+$. It reflects a razor's-edge view of a 0 recommendation. The axiom "pushes" the systems towards strict recommendations. (Without such an axiom, systems may almost always recommend 0.)

AXIOM 2. (Positive response) *Let $w \notin N$, $s \in \bar{V}$, $G = (N, V_+, V_-, E)$, and $G' = (N \cup \{w\}, V_+ \cup \{w\}, V_-, E \cup \{(s, w)\})$. If $s \notin R_-(G)$ then $s \in R_+(G')$.*

Note that the above axiom is presented asymmetrically in terms of \pm votes and recommendations. In combination with the Symmetry axiom, the corresponding version with $-$ votes and recommendations follows directly. We use an asymmetric presentation for readability in several of the axioms.

The next axiom, Independence of Irrelevant Stuff (IIS) captures the semantics of recommendation systems discussed in the introduction: a source node is viewed as consulting neighboring agents in the trust graph, who consult their neighbors etc., while agents who have formed opinions just vote according to their opinion. This means that when considering the recommendation for a particular source node in a particular trust graph, where part of the agents vote (perhaps based on first-hand experience), feedback from these agents is independent of who they trust (i.e., they trust themselves infinitely more than others) and the recommendation system should consider only reachable nodes and ignore links out of voters. While one may consider other types of semantics, something similar to this axiom appears in many previously designed systems.

AXIOM 3. (IIS) *Let $G = (N, V_+, V_-, E)$ and $e \in V \times N$ be an edge leaving a voter. Then for the subgraph $G' =$*

$(N, V_+, V_-, E \setminus \{e\})$ in which e has been removed, $R(G) = R(G')$. Similarly, if $v \in N$ is a node not reachable from $s \in \bar{V}$, then for the subgraph G'' in which node v (and its associated edges) have been removed, $R(G, s) = R(G'', s)$.

When we write $R(G) = R(G')$, as in the above, it means that the recommendations on the two voting networks are identical.

The following requirement deals with some minimal rationality we wish to attribute to the agents; as in the classical theory of choice we are willing to assume something about the vote of an agent who has no a priori opinion only in extreme cases. The neighborhood-consensus axiom does just that: if all the agents trusted by a node v vote +, and no other nodes touch v 's neighbors, then the recommendation of any other nonvoter u will remain unchanged if v instead becomes a + voter.

AXIOM 4. (Neighborhood Consensus) Fix a voting network $G = (N, V_+, V_-, E)$, and let $s, u \in \bar{V}$ be distinct nonvoters. Suppose u has at least one outgoing edge, and suppose that each outgoing edge $(u, v) \in E$ points to v such that $v \in V_+$ and v has no (incoming or outgoing) neighbors other than u . Let $G' = (N, V_+ \cup \{u\}, V_-, E)$. Then $R(G, s) = R(G', s)$.

4.1 Transitivity

Transitivity is a central concept in the axiomatization of voting [6]. In our context, we consider the case where the underlying trust graph is fixed, but the system needs to deal with more than one item, and different subsets of nodes vote on different items. The definition of transitivity is based on the following binary relation, which describes whether a specified source node trusts a set of nodes A more than another set of nodes B .

DEFINITION 3. Let $G = (N, V_+, V_-, E)$ be a voting network. If $s \in R_+(G)$, then we say that s trusts V_+ more than V_- relative to multigraph (N, E) .

For each source node, a binary relation is generated. The transitivity axiom stipulates that each of these binary relations must be transitive.

AXIOM 5. (Transitivity) For all multigraphs (N, E) , $s \in \bar{V}$, and disjoint $A, B, C \subseteq N$, if, relative to (N, E) , s trusts A more than B and s trusts B more than C , then s trusts A more than C .

4.2 Trust Propagation

In this section, we consider *propagation of trust*. Intuitively, if u trusts v and v trusts w , then u should trust w , at least to some extent. Much has been written about trust propagation within social networks (see, e.g., [12]) and the axiom below is a conservative interpretation that agrees with much of the literature.

One would like to say that if u trusts nonvoter v , and v trusts w , then we can simply add an edge from u to w without changing anything. However, our system is supposed to reflect degrees of trust, and this would falsely inflate such trust. Instead, we count edges as follows. Suppose there are k edges leaving v (that don't point to u). Suppose that there happen to be k edges from u to v . Then we can remove k edges from u to v and replace them by k new edges from u to the k nodes that v trusts (besides u), and no recommendations should change.

AXIOM 6. (Trust propagation) Let voting network $G = (N, V_+, V_-, E)$, $u \neq v \in \bar{V}$, and suppose that the edges leaving v (besides those to u) are $(v, w_1), \dots, (v, w_k)$, ($w_i \neq u$) for some integer k . Suppose that E contains exactly k copies of (u, v) . Then, for $E' = (E \uplus \{(u, w_1), \dots, (u, w_k)\}) \setminus \{(u, v) * k\}$ and $G' = (N, V_+, V_-, E')$, we have that $R(G) = R(G')$.

Another natural axiom is *scale invariance*. Loosely speaking, this means that the amount of trust placed in a node is relative.

AXIOM 7. (Scale invariance) In voting network $G = (N, V_+, V_-, E)$, $u \in \bar{V}$, and $k \geq 1$. Let $G' = (N, V_+, V_-, E \uplus E')$, where E' is the multiset containing k copies of each of the edges leaving v . Then $R(G) = R(G')$.

It states that we can duplicate all edges leaving a node an arbitrary number of times without changing recommendations.

4.3 Majority and Groupthink

The majority axiom states that a node's recommendation should equal the majority of the recommendations and votes of its trusted neighbors. We say the sign of an edge is *positive*, *negative*, or *neutral* if it points to a node with positive vote or recommendation, negative vote or recommendation, or 0 recommendation, respectively.

AXIOM 8. (Majority) Let $G = (N, V_+, V_-, E)$ be a voting network, and let $s \in \bar{V}$ be any nonvoter. Then, the recommendation of s should be equal to the majority of the signs of the edges leaving s .

We are using the strict notion of majority as defined in Section 2. This choice is somewhat arbitrary, though it fits well with the next axiom. Also note that one can further axiomatize majority itself, but we leave it as is for the sake of brevity. Unfortunately, the majority axiom by itself does not imply unique recommendations on cyclic graphs (think about a graph with two nonvoters that point to each other). Instead, we consider the following property.

Groupthink refers to a social phenomenon in which an entire group of people arrive at a ridiculous conclusion through self-reinforcing interactions. The no-groupthink axiom rules this out and imposes strong semantics on the system. There are two parts to this axiom. First, we consider the case that an entire group of nonvoters are all recommended +. This strong position should be based on something external, since no member voted. The requirement is that, among the edges leaving the group, a majority must point to nodes with + votes or recommendations. Furthermore, if a majority of the edges leaving the group point to nodes with + votes or recommendations, then the group must contain at least one node with + recommendation.

Since no-groupthink is inconsistent for general directed graphs, we define it for specific graphs. We say that a recommendation system R avoids groupthink for G if the following holds:

AXIOM 9. (No groupthink) Let $S \subseteq \bar{V}$ be a nonempty set of nonvoters. Let E' be the multiset of edges in E from S to $N \setminus S$. (a) If $S \subseteq R_+(G)$, then a strict majority of the edges in E' must point to nodes in $R_+(G) \cup V_+$. (b) If a strict majority of the edges in E' point to nodes in $R_+(G) \cup V_+$, then $S \cap R_+(G) \neq \emptyset$.




Figure 3: (i) Example demonstrating the impossibility of axioms 1-5. (ii) In this example, s trusts $\{u\}$ more than $\{v\}$, $\{v\}$ more than $\{w\}$ and $\{w\}$ more than $\{u\}$. The details are described in the text.

When S is a singleton the no-groupthink axiom reduces to the majority axiom. Hence no-groupthink can be interpreted as a generalization of the majority axiom to larger sets. Note that the no-groupthink axiom is incompatible with the scale-invariance axiom.

5. ANALYSIS SKETCH

In this section, we give intuition underlying the proofs of the three theorems from Section 1.2. Formal proofs are deferred to the Appendix.

For Theorem 1, the example in Figure 3(i) is used to demonstrate impossibility. Relative to the depicted graph, one can show that s trusts set A more than B , B more than C , and C more than A , contradicting transitivity. To show that the set of axioms is a minimal impossible set, we construct five recommendation systems, each one satisfying all but one of the axioms. For example, the RW recommendation system satisfies axioms 1-4 but does not satisfy axiom 5 (transitivity).

Figure 3(ii) gives another example worthy of discussion. When u votes $+$, v votes $-$, and w does not vote, it is reasonable that s should be $+$ due to its extra path to u through w and the fact that multiple edges in our system reflect additional trust. However, continuing along these lines, symmetry leads to a cycle in the transitivity relation. While this argument does not follow from axioms 1-5, it is suggestive of problems due to transitivity.

For Theorem 2, we first apply a sequence of transformations to the graph, to get the graph to a point in which all edges lead to voters. Each of these transformations can be shown, by the axioms, not to change any recommendations. Moreover, none of these transformations change the recommendations computed by RW. Thus, the problem is reduced to showing that the theorem holds when all edges point to voters.

The transformation process iterates through the nonvoters, and, for each non-voter, it removes all incoming edges by a combination of edge duplication and trust propagation. When we are left with edges only to voters, axioms 1-4 are still not sufficient to determine the recommendations of nodes due to possible multiple edges. However, by applying the trust propagation axiom again (in some sense in reverse), we complete the argument.

For Theorem 4, part (a) on rooted trees is much simpler than part (b). For part (a), the no-groupthink axiom on

an individual implies directly that its vote will equal the majority of its outgoing neighbors' recommendations. This can be applied from the sinks upwards to get uniquely the MoM system.

For part (b), our proof is by contradiction. We begin by assuming that the recommendations do not match those of the min-cut system and use a case analysis to show that, in each case, some set of nodes engages in groupthink.

6. INCENTIVE COMPATIBILITY

Incentive compatibility is a desirable property for voting systems. In our context, it not only means that an agent will have incentive to vote according to its prior opinion, but it means that an agent (or a set of agents) cannot be more effective by creating false edges (or even false nodes).

To discuss incentive compatibility, one needs to discuss preferences. We say an agent has *positive preferences* if an agent prefers a set of recommendations that is, node by node, larger than or equal to another set of recommendations (where $+$ is larger than 0 and $-$, and 0 is larger than $-$). Similarly, for *negative preferences*.

We say a recommendation system is *incentive compatible* if (1) every set of agents with positive preferences is at least as happy with the outcomes when they all vote $+$ and report trust honestly, as they would be in any other scenario that they can collude to create; and (2) the same is true for every set of agents with negative preferences when they all vote $-$.

For example, suppose a set of agents T have positive preferences, and node $s \notin T$ receives a 0 or $-$ when all agents in S vote $+$. Incentive compatibility implies that s cannot receive a greater recommendation when agents in T take any of the following strategic actions (sometimes called *Sybil attacks*):

- T creates an arbitrarily large set of fictitious agents F .
- T and F create arbitrary edges between themselves, and arbitrary outgoing edges from themselves to other agents. (They are not allowed to alter edges which start at non-malicious agents.)
- Arbitrary subsets of $S \cup F$ vote $+$, vote $-$, and do not vote.

This implies that the set S will never strictly prefer any outcome (under their control) other than that achieved when they all vote $+$. It is relatively straightforward to see that RW satisfies incentive compatibility, as the best strategy for a group of agents to maximize another node's vote is to maximize the probability that the first voting node encountered votes $+$. Before a random walk reaches that malicious agents, it is out of their control. Once it reaches them, they might as well terminate the walk with a $+$ vote.

It is similarly straightforward to verify that the MoM system also satisfies incentive compatibility if we enforce the fact that the cheating agents cannot create any cycles in the graph. Lastly, the min-cut system also obeys the above type of incentive compatibility. Note that personalized PageRank does not satisfy incentive compatibility, since a voting node with positive preferences might have an incentive to alter its outlinks.

Incentive compatibility is closely linked to axiom 3 (IIS), but they are not equivalent.

7. CONCLUSIONS

In this paper, we initiated the axiomatic study of trust-based recommendation systems. This allows for rigorous evaluation of recommendation systems. Our work deals both with a normative approach, where the ramifications of natural postulates are considered, and with the descriptive approach where we aim at fully characterizing particular systems. In particular, we have found five basic axioms that cannot be satisfied simultaneously, for which any proper subset can be satisfied. In addition, we have given a sharp characterization of the random walk recommendation system. This was obtained by replacing an axiom capturing a notion of transitivity with ones capturing trust propagation and duplication.

Acknowledgments. We thank the attendees of the “DI-MACS Workshop on the Boundary between Economic Theory and Computer Science” for their many helpful comments.

8. REFERENCES

- [1] A. Altman and M. Tennenholtz. On the axiomatic foundations of ranking systems. In *Proc. 19th International Joint Conference on Artificial Intelligence*, pages 917–922, 2005.
- [2] A. Altman and M. Tennenholtz. Ranking systems: the pagerank axioms. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 1–8, New York, NY, USA, 2005. ACM Press.
- [3] A. Altman and M. Tennenholtz. An axiomatic approach to personalized ranking systems. In *Proc. 20th International Joint Conference on Artificial Intelligence*, 2006.
- [4] A. Altman and M. Tennenholtz. Quantifying incentive compatibility of ranking systems. In *Proc. of AAAI-06*, 2006.
- [5] A. Altman and M. Tennenholtz. Incentive compatible ranking systems. In *Proceedings of the 2007 International Conference on Autonomous Agents and Multiagent Systems (AAMAS-07)*, 2007.
- [6] K. Arrow. *Social Choice and Individual Values (2nd Ed.)*. Yale University Press, 1963.
- [7] P. Avesani, P. Massa, and R. Tiella. A trust-enhanced recommender system application: Moleskiing. In *SAC '05: Proceedings of the 2005 ACM symposium on Applied computing*, pages 1589–1593, New York, NY, USA, 2005. ACM Press.
- [8] Y. Bakos and C. N. Dellarocas. Cooperation without enforcement? a comparative analysis of litigation and online reputation as quality assurance mechanisms. MIT Sloan School of Management Working Paper No. 4295-03, 2003.
- [9] A. Borodin, G. O. Roberts, J. S. Rosenthal, and P. Tsaparas. Link analysis ranking: algorithms, theory, and experiments. *ACM Trans. Inter. Tech.*, 5(1):231–297, 2005.
- [10] A. Cheng and E. Friedman. Sybilproof reputation mechanisms. In *P2PECON '05: Proceeding of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, pages 128–132, New York, NY, USA, 2005. ACM Press.
- [11] R. Dash, S. Ramchurn, and N. Jennings. Trust-based mechanism design. In *Proceedings of the Third International Joint Conference on Autonomous Agents and MultiAgent Systems*, pages 748–755, 2004.
- [12] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 403–412, 2004.
- [13] T. Haveliwala. Topic-sensitive pagerank. In *WWW '02: Proceedings of the 11th international conference on World Wide Web*, pages 517–526, 2002.
- [14] T. Haveliwala, S. Kamvar, and G. Jeh. An analytical comparison of approaches to personalizing pagerank, 2003. Technical report, Stanford University.
- [15] D. Houser and J. Wooders. Reputation in auctions: Theory and evidence from ebay. Working Paper, University of Arizona, 2000.
- [16] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.
- [17] R. Levien. *Attack Resistant Trust Metrics*. PhD thesis, University of California, Berkeley, 2002.
- [18] P. Massa and P. Avesani. Controversial users demand local trust metrics: An experimental study on epinions.com community. In *Proc. of AAAI-05*, pages 121–126, 2005.
- [19] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report, Stanford University, 1998.
- [20] Pennock D. M., Horvitz E., and Giles C. L. Social Choice Theory and Recommender Systems: Analysis of the Axiomatic Foundations of Collaborative Filtering. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-2000)*, pages 729–734. AAAI Press, 2000.
- [21] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay’s reputation system. Working Paper for the NBER workshop on empirical studies of electronic commerce, 2001.
- [22] P. Resnick, R. Zeckhauser, R. Friedman, and E. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.
- [23] A. P. Singh, A. Gunawardana, C. Meek, and A. C. Surendran. Recommendations using absorbing random walks. Manuscript under submission.
- [24] G. Slutzki and O. Volij. Ranking participants in generalized tournaments. *International Journal of Game Theory*, 33(2):255–270, 2005.
- [25] M. Tennenholtz. Reputation systems: An axiomatic approach. In *Proceedings of the 20th conference on uncertainty in Artificial Intelligence (UAI-04)*, 2004.

APPENDIX

A. PROOF OF THEOREM 1

First, we show that there exists a recommendation that satisfies any proper subset of the axioms. In the following cases, we consider a voting network $G = (N, V_+, V_-, E)$ and source $s \in N$.

1. System without symmetry. Consider the trivial recommendation system that always recommends $+$. This system trivially satisfies the remaining axioms.

2. System without positive response. Consider the

trivial recommendation system that always recommends 0. This system trivially satisfies the remaining axioms. (Recall that axiom 4, neighborhood consensus, refers only to distinct non-voters.)

3. System without neighborhood-consensus. Consider the recommendation system which assigns to each non-voter u the recommendation equal to $\text{sgn}(a - b)$ where a is the number of edges from u to positive voters and b is the number of edges to negative voters. It is symmetric, neutral, and obeys positive response. It obeys IIS since immediate neighbors of s are reachable and can not be removed. Transitivity is implied by the fact that this “local majority” system simply counts the number of edges leading to a set of nodes; since the number of edges connecting to $X \subseteq \text{Succ}(s)$ is fixed, it obeys transitivity.

4. A System Without IIS. In order to design a recommendation system that satisfies all axioms except IIS, define S_+ to be the set of nodes $u \in N$ that have at least one outgoing neighbor, and all of whose outgoing neighbors are in V_+ and have no neighbors other than u . (These are essentially the nodes described in the consensus axiom, except that it applies to voters as well, e.g., $S_+ \cap V_-$ is not necessarily empty.) Similarly define S_- . For any nonvoter $s \in \bar{V}$, the recommendation to s is $+$ if $s \in S_+$, $-$ if $s \in S_-$. For all of the remaining nonvoters, the recommendation is $\text{sgn}(|S_+ \cup V_+| - |S_- \cup V_-|)$.

It is easy to see that this recommendation system satisfies symmetry. The positive response axiom follows from the fact that the recommendation 0 for s can be generated only when $|S_+ \cup V_+| = |S_- \cup V_-|$. Adding a new positive voting node w that is pointed to by only s must increase $|S_+ \cup V_+|$ and cannot increase $|S_- \cup V_-|$, since $s \notin S_- \cup S_+ \cup V_- \cup V_+$. To see that the system satisfies neighborhood consensus, note that moving a nonvoter from S_+ to V_+ doesn’t change any the sets in any other way. Finally, to see that the system satisfies transitivity, consider disjoint $A, B, C \subseteq N$ and $s \in N$ such that s trusts A more than B and B more than C . Note that S_+ depends only on V_+ and is independent of V_- . Also note that switching the sign of the votes switches the sets S_+ and S_- due to the symmetry of the algorithm.

Next consider two cases. In case 1, suppose that $s \in S_+$ when $V_+ = A$. This means that s will receive a positive recommendation when $V_- = C$ so s trusts A more than C . In case 2, $s \notin S_+$ when $V_+ = A$. Since s trusts A more than B and B more than C , this means in particular that $s \notin S_+$ when $V_+ = B$ nor when $V_+ = C$ (and $s \notin S_-$ when $V_- = B$ or $V_- = C$). Hence, in all cases, s ’s vote will equal $\text{sgn}(|S_+ \cup V_+| - |S_- \cup V_-|)$. Thus $|S_+ \cup V_+|$ is a measure of trust of s in a set V_+ , and hence the system obeys transitivity.

5. A System Without Transitivity. Theorem 2 implies that RW, the random walk recommendation system, satisfies axioms 1-4.

6. Inconsistency Now we prove that there is no recommendation system satisfying axioms 1-5. Consider the graph (V, E) depicted in Figure 3(i). We claim that s trusts A more than B and s trusts B more than C , but s does not trust A more than C . This will prove impossibility.

To see that s trusts A more than B , consider the voting network $G = (V, A, B, E)$. By the IIS axiom, we can ignore the edges from B to C without changing the recommendation to s , to get voting network G' . Next imagine removing one of the nodes in A (and its associated edge) to get voting

network G'' . Since G'' is a star graph with three $+$ votes, three $-$ votes, and 2 nonvoters aside from s , the recommendation for s has to be 0 by the symmetry axiom. By positive response, therefore, the recommendation for s in G' has to be $+$ and hence s trusts A more than B .

To see that s trusts B more than C , consider the voting network $G = (V, B, C, E)$. Again by the IIS axiom, we can ignore the edges from B to C without changing the recommendation to s . We are now again in a star graph with 3 $+$ votes, 2 $-$ votes, and 4 0 votes. By symmetry and positive response, as above, the recommendation to s must be $+$, which is what we needed.

Finally, to see that s does not trust A more than C , consider the voting network $G = (V, A, C, E)$. By three applications of neighborhood consensus we can take the votes of B to be all $-$. By IIS we can transform the graph yet again into a star graph, this time with 5 $-$ votes and 4 $+$ votes. As reasoned above, the recommendation to s must be $-$, and hence s does not trust A more than C . \square

B. PROOF OF THEOREM 2

We first apply a sequence of transformations to the graph, to get the graph to a point in which all edges lead to voters or sink nonvoters. Each of these transformations can be shown, by the axioms, not to change any recommendations. Moreover, none of these transformations change the recommendations computed by RW. Hence, it suffices to show the theorem for the case where all edges point to voters or sink nonvoters.

The transformation removes all edges pointing to any single nonvoter (that is not a sink), as follows. Take a node v which is a nonvoter but has out-degree $k > 0$. Take any other node u that points to v $\ell > 0$ times. We will apply trust propagation to remove these edges to v , without adding other edges to v . To do this, we apply ℓ -fold edge duplication (scale invariance) to v and k -fold edge duplication (scale invariance) to u without changing any recommendations. This makes the number of edges from u to v equal to ℓk and the out-degree of v also equal to ℓk . Hence, we can apply trust propagation to remove all edges from u to v . This can be applied in turn to each node that points to v , to get a graph where all edges point to voters or sink nonvoters.

Note that each of these transformations does not change the recommendations of the RW system, either. The reason is that RW obeys scale invariance and trust propagation, which can be seen easily from the random walk definition. When the random walk reaches a non-sink nonvoter, it continues on a random edge out of that nonvoter. The trust propagation simply provides a “shortcut” which exactly simulates two steps of the random walk at this point.

So, without loss of generality, it suffices to show the RW system is correct for a graph with only edges pointing to voters and sink nonvoters. Consider any source s and the subgraph reachable from that source via directed edges through nonvoters. If no voter had multiple edges pointing to it, then the theorem would follow from positive response and symmetry. The reason is that when there were an equal number of $+$ and $-$ voters, by symmetry the recommendation of s would have to be 0, regardless of the structure of edges to nonvoter sinks. By positive response, adding any number of edges from s to more new $+$ voters would cause the recommendation of s to become positive. Similarly, if there were

more (nonparallel) edges to $-$ voters.

The case that remains is where there are multiple parallel edges to voters. Say we have a $+$ -voter v with k incoming edges (and, without loss of generality, no outgoing edges). By neighborhood consensus, we can change such a voter to a nonvoter and add a number of outgoing edges to k new, distinct $+$ voters that have no other neighbors. Now we can apply trust propagation to v to get a new graph where all edges pointing to v have been replaced by edges pointing directly to new nonvoters, without changing recommendations or the behavior of RW. In addition, we can remove v without changing s 's recommendation, since there is no path from s to v . We can apply this procedure to each $+$ and $-$ voter with more than one incoming edge, until all voters have in-degree 1. The result for this case then follows from the argument in the above paragraph.

C. PROOF OF THEOREM 3

Without loss of generality, we assume there are no edges between pairs of voters. Part (a) follows trivially from the majority axiom and the fact that any directed acyclic graph can be partitioned into levels, where each edge is from a node in a higher level to a node in a lower level.

For Theorem 4(b), we first state an interesting property of the min-cut system.

LEMMA 1. *Let $G = (N, V_+, V_-, E)$ be a voting network and $R(G)$ be the recommendations of the min-cut system. Let C the multiset of edges between nodes in $R_+(G)$ and nodes in $V_- \cup R_-(G) \cup R_0(G)$ plus with those between nodes in $R_-(G) \cup R_0(G)$ and nodes in V_+ . Then C is a min-cut.*

Note that this lemma is asymmetric, we have essentially put all 0 recommendations with the negative recommendations.

PROOF. Given cuts C and C' define the *max* of the two min-cuts as follows. For any cut C , let the node-set C_+ be the set consisting of $+$ -voters and also nonvoters that can reach any $+$ voter after C has been removed. Given min-cuts C and C' , and their corresponding node sets C_+ and C'_+ , the *max* of the two min-cuts is defined to be the cut that cuts any edges between nodes in $C_+ \cap C'_+$ and $N \setminus (C_+ \cap C'_+)$. (Note that the node-set corresponding to the max of the two cuts is simply $C_+ \cap C'_+$.) By repeated application of this observation, we get the lemma, because the intersection of all node-sets corresponding to min-cuts gives exactly the cut described by the lemma.

It remains to show that the max of C and C' is a min-cut by itself. To do this, it is helpful to define the *min* of two min-cuts, which is the completely symmetric notion to the max. Let the max be A and the min be B . It is not hard to see that multisets $A \uplus B \subseteq C \uplus C'$, where \uplus denotes multiset union. The reason is that the number of times a given edge is cut by C and C' (0, 1, or 2) is at least as great as the number of times it is cut by A and B (those edges between $C_+ \setminus C'_+$ and $C'_+ \setminus C_+$ are not cut at all by A or B). Thus $|A| + |B| \leq |C| + |C'|$. Furthermore, A and B are both cuts. Hence, at least one of them, say A , must no larger than C and C' , which have the same size. But $|A| = |C| = |C'|$ since C and C' are min-cuts. Hence, $|B| = |A|$ and they are all min-cuts. \square

We first argue that the min-cut system satisfies the groupthink axiom. Consider first part (a) of groupthink. Suppose

$S \subseteq R_+(G)$ is a subset of the nonvoters that are all given positive recommendations. These nodes are all connected to positive voters in *every* min-cut. Take the min-cut defined in the above lemma. Now, for the purposes of contradiction, suppose S does not have a majority of positive edges leaving it. Then we claim that we could find another min-cut C' in which no node in S is connected to a $+$ voter, which violates the definition of the min-cut system. In particular, we would take C' to be the cut which is equal to C , except that it also cuts all positive edges leaving S and does not cut any negative or neutral edges leaving S . This is a cut because C was a cut and no new paths between positive and negative voters have been created – they would have to go through S but S has no positive edges and thus is disconnected from the positive voters. Furthermore, C' would not be larger than C because we have assumed that S does not have a majority of positive edges leaving it. Hence, we have a min-cut and the desired contradiction. The argument for the case where S all has negative recommendations follows by symmetry of the system.

We next argue that the min-cut system satisfies Groupthink (b). For the purposes of contradiction, suppose we have a set $S \subseteq R_-(G) \cup R_0(G)$ such that S has a majority of positive outgoing edges. (The other case, where $S \subseteq R_+(G) \cup R_0(G)$ such that S has a majority of negative outgoing edges, is similar.) Again, take the min-cut C defined by the above Lemma. Consider the cut C' which is the same as C but where we have also cut all negative and neutral edges from S and haven't cut any positive edges. As in the previous argument, this remains a cut and is smaller than C , giving us the desired contradiction.

It remains to show that the min-cut system is the unique system implied by the groupthink axiom. Suppose not. Suppose that on some undirected graph, different recommendations $R'(G)$ satisfy the groupthink axiom.

Case 1: The set $S = (R'_-(G) \cup R'_0(G)) \cap R_+(G)$ is nonempty. Since we have shown that the min-cut system satisfies groupthink and $S \subseteq R_+(G)$, this means that a strict majority of edges between S and $N \setminus S$ are between a node in S and a node in $V_+ \cup (R_+(G) \setminus S)$. However, by definition of S , $(R_+(G) \setminus S) \cap (R'_-(G) \cup R'_0(G))$ is the empty set, or in other words $R_+(G) \setminus S \subseteq R'_+(G)$ since every nonvoter has either a $-$, 0, or $+$ recommendation in R' . This gives us the desired contradiction because we have more than half of the edges leaving S pointing to nodes in $R'_+(G) \cup V_+$, which contradicts Groupthink (a).

Case 2: The set $(R'_+(G) \cup R'_0(G)) \cap R_-(G) \neq \emptyset$. This follows from the previous case by symmetry.

Case 3: The sets $(R'_-(G) \cup R'_0(G)) \cap R_+(G) = (R'_+(G) \cup R'_0(G)) \cap R_-(G) = \emptyset$, but $T = R'_+(G) \cap R_0(G) \neq \emptyset$. Again, because the min-cut system satisfies groupthink, the set T which is unanimously 0 under R cannot have a strict majority of its edges pointing to nodes in $(R_+(G) \cup V_+) \setminus T$. However, it is unanimously positive under R' , so it must have a strict majority of edges pointing to nodes in $(R'_+(G) \cup V_+) \setminus T$. However, by our choice of T , these two sets are equal.

Case 4: The sets $(R'_-(G) \cup R'_0(G)) \cap R_+(G) = (R'_+(G) \cup R'_0(G)) \cap R_-(G) = \emptyset$, but $R'_-(G) \cap R_0(G) \neq \emptyset$. This follows from the previous case by symmetry.

The above four cases cover all possibilities in which $R' \neq R$. \square