False-Name-Proof Recommendations in Social Networks

Markus Brill
Duke University
brill@cs.duke.edu

Vincent Conitzer
Duke University
conitzer@cs.duke.edu

Rupert Freeman
Duke University
rupert@cs.duke.edu

Nisarg Shah
Carnegie Mellon University
nkshah@cs.cmu.edu

Abstract

We study the problem of finding a recommendation for an uninformed user in a social network by weighting and aggregating the opinions offered by the informed users in the network. In social networks, an informed user may try to manipulate the recommendation by performing a false-name manipulation, wherein the user submits multiple opinions through fake accounts. To that end, we impose a *no harm axiom*: false-name manipulations by a user should not reduce the weight of other users in the network. We show that this axiom has deep connections to false-name-proofness. While it is impossible to design a mechanism that is best for every network subject to this axiom, we propose an intuitive mechanism LEGIT$^+$, and show that it is uniquely optimized for small networks. Using real-world datasets, we show that our mechanism performs very well compared to two baseline mechanisms in a number of metrics, even on large networks.

1 Introduction

Consider the following problem. An agent wants to receive a recommendation on a specific item—say, a movie the agent has not previously watched. Others have evaluated this item, perhaps by giving it a "thumbs up" or "thumbs down" (0 or 1), or by rating it on a more detailed scale, say, from 0 to 5. We want to give the agent in question an aggregate rating, such as "73% positive" or 2.7. Alternatively, perhaps the question is merely whether to recommend this item to the agent at all, in which case the aggregate outcome must be binary. How should we arrive at this aggregate outcome?

For simplicity, let us assume that we do not have any information about which agents have preferences most similar to the agent in question. In this case, a natural approach is to simply take the average of all the ratings so far. One problem is that if ratings are not binary, this is not strategy-proof: when the current average is 2.7, an agent who feels the item is a 4 may prefer to report 5 to pull the average closer to his evaluation. As is well known in social choice theory, a good alternative is to choose the *median* rating instead: this is in fact group-strategy-proof (when preferences are single-peaked, as is likely to be the case here) (Moulin, 1980; Barbera, Sonnenschein, and Zhou, 1991; Border and Jordan, 1983). Note that for binary ratings, the median is simply the majority choice.

The median, however, remains vulnerable to another type of manipulation, commonly known as *false-name manipulation* (Yokoo, Sakurai, and Matsubara, 2004): an agent can rate the same item many times by opening fake accounts, and move the median closer to his evaluation. Thus, the median is not false-name-proof. In fact, without imposing further structure on the problem, no reasonable rule is false-name-proof (Conitzer, 2008; Todo, Iwasaki, and Yokoo, 2010, 2011). On the other hand, if we assume that agents are organized as the nodes of a (say, undirected) social network, possibilities open up (Conitzer et al., 2010). For example, rather than reporting to the agent the median of all the ratings, we can simply report the median of his friends' ratings. Assuming that the agent will not be duped into befriending fake accounts, this will in fact be false-name-proof.

The downside of this methodology is that, with the exception of very popular items, none or very few of the agent's friends may have rated the item. Consequently, the median-of-friends rule conveys too little information. Could we include the friends of the agent's friends as well? Done naïvely, this may give the friends an incentive to create many fake friends of their own. But more subtly, when a friend does *not* rate the item, we can pretend that his rating was the median of *his* friends' ratings. This does not give the friends an incentive to create fake accounts: all this would do for them is change their own hallucinated ratings, but they can more easily just specify those ratings directly. This median-of-medians approach closely resembles the majority-of-majorities rule from Andersen et al. (Andersen et al., 2008). Note, however, that this rule ends up double-counting the rating of an agent who is a friend of two of the agent's friends. Can we circumwent this issue? Also, can we retrieve ratings from deeper in the social network?

Our results. In this paper, we focus on a two-step approach. In the first step, we use a *weight-selecting mechanism* to assign weights to the agents offering an opinion/rating, called *voters*, without looking at their opinions (thus, looking only at the network structure). In the second step, we perform a weighted aggregation of the opinions to output a recommendation by only looking at the weights assigned to voters and their opinions. To make the weight-selecting mechanism robust to false-name manipulations, we impose a *no harm axiom*: false-name manipulations by an agent should not reduce the weight of other agents in the network.

We show that with weighted median aggregation, the no

harm axiom implies false-name-proofness (Theorem 1) and, under some conditions, is actually equivalent to false-name-proofness (Theorem 2). We thus focus on designing weight-selecting mechanisms subject to this axiom.

We focus on the case where, ideally, we would like to weight the voters uniformly. As explained in detail in Section 2, this is for multiple reasons. While this does not utilize the network structure for inferring the closeness in opinions of two nodes, it clearly outlines how to use the network structure for a distinct purpose — achieving the no harm axiom. Section 6 discusses how our results can be extended to take into account correlation among opinions. Second, weighting the voters equally can indeed be ideal, e.g., when aggregating independent noisy estimates of an underlying objective ground truth (see Section 5), or when the goal is not to find a recommendation but to conduct a fair vote.

Unfortunately, weighting all the voters uniformly violates the no harm axiom. What is the "most uniform" weight vector we can return subject to this axiom? In order to formalize what "more uniform" means, we use the classic leximin criterion that compares weight vectors by their smallest weights (preferring the vector with greater smallest weight), and then breaks ties using the second smallest weights, and so on. We show that a weight-selecting mechanism cannot always return the leximin-optimal weight vector subject to the no harm axiom (Theorem 4). We then present an intuitive mechanism and show that it is uniquely optimized for small networks subject to the no harm axiom (Theorem 5), that is, (informally) if a mechanism outputs a more uniform weight vector than our mechanism does on some network, then there is a *strict subgraph* of the network on which our mechanism outputs a more uniform weight vector.

Using a non-trivial result from graph theory (Hopcroft and Tarjan, 1973), we show that our mechanism can be computed in linear time in the size of the network (Theorem 6). In Section 5, we present experiments with real-world social networks in which our mechanism significantly outperforms two baseline mechanisms in a number of metrics.

Related work. Recommendation systems have been studied extensively in the machine learning literature, see, e.g., (He and Chu, 2010; Adomavicius and Tuzhilin, 2005; Ricci, Rokach, and Shapira, 2011; Bennett and Lanning, 2007). Popular techniques include content-based recommendation (Pazzani and Billsus, 2007), where the decision of whether to suggest an item to a target user is made by considering the attributes of the item and the target user's previously expressed preferences; collaborative filtering (Golbeck, 2006), where the preferences of other users in the network are given, and their similarity with the target user's preferences is learned to find a good recommendation; or, both combined (Balabanović and Shoham, 1997). In contrast, we solely focus on the use of the social network structure to design recommendation mechanisms that are robust to false-name manipulations.

Besides the works cited previously, false-name manipulations have also been studied rigorously in a variety of anonymous environments, such as combinatorial auctions (Yokoo, Sakurai, and Matsubara, 2001; Yokoo, 2003; Yokoo, Sakurai, and Matsubara, 2004; Todo et al., 2009; Iwasaki et al., 2010), matching (Todo and Conitzer, 2013), and voting (Wagman and Conitzer, 2008).

2 Model

We are given a *social network* (or simply, a network), which is an undirected simple graph[1] denoted G. The set of nodes and the set of edges of G are denoted $V(G)$ and $E(G)$ (or V and E, when the graph is clear from the context), respectively. For $T \subseteq V$, let G_T denote the subgraph of G induced by T.

Our task is to find a recommendation for a given node $v^* \in V$. This task could arise in a number of contexts: we may want to decide whether to recommend a given movie or restaurant to an individual (in which case, we want a binary recommendation), or we may want to show the rating of the movie or restaurant (in which case, we no longer want a binary recommendation). To aid our decision-making, a set of nodes $S \subseteq V \setminus \{v^*\}$ offer their personal opinion. We call these nodes *voters*, and denote the opinion offered by voter $v \in S$ as r_v. Target node v^* is not a voter itself. As we explain in Section 3, the mechanisms of our interest must discard voters not connected to v^*; thus, for simplicity we assume that G is connected and has at least one voter.

Weight-Selecting Mechanisms: In this paper, we are interested in finding recommendations through a two-step approach: i) using a *weight-selecting mechanism* to assign a weight to each voter in the network only as a function of the network structure G, the subset of voters S, and the target node v^* (thus, *independent* of the voters' opinions), and ii) using a *weighted aggregation function* that takes as input the weights assigned by the weight-selecting mechanism and the voters' opinions, and outputs the final recommendation. Popular choices for the weighted aggregation function include weighted mean and weighted median; Section 3.1 discusses how this choice impacts the overall recommendation system. For the remaining parts of the paper, we are only interested in studying weight-selecting mechanisms and their properties (the first step). For a weight-selecting mechanism — ignorant of the voters' opinions — a problem instance is given by the tuple (G, S, v^*).

Definition 1 (Weight-Selecting Mechanisms)**.** *Given an instance (G, S, v^*), a weight-selecting mechanism outputs a weight vector $\mathbf{w} = (w_v)_{v \in S}$ such that $w_v \geq 0$ for $v \in S$, and $\sum_{v \in S} w_v = 1$.*

Weight-selecting mechanisms are compelling because they allow harmonious aggregation of opinions of various formats, ranging from binary to real-valued opinions.

False-Name Manipulations: In the absence of additional restrictions, one can simply choose the weight-selecting mechanism that returns the most appropriate weight vector for the setting of interest. In this paper, however, we consider an important restriction that stems from game-theoretic considerations: *preventing false-name manipulations*.

[1] A simple graph has no self-loops and at most one edge between every pair of vertices.

Online social networks typically lack a proof of authenticity of nodes, thus allowing users to easily create fake accounts. In this case, a weight-selecting mechanism may inadvertently provide an incentive to a malicious user for creating multiple fake accounts and voting through them in order to gain a higher total weight, and thus a greater influence on the final recommendation. Such manipulations are known as *false-name manipulations* or *sybil attacks*.

In a false-name manipulation, the malicious node in the network can easily create any desired subset of edges among the identities it controls: its own node, and the fake nodes it creates. Altering edges with other real nodes (e.g., creating new edges or deleting existing edges), on the other hand, is often more costly. Given that (arguably) recommendations are not the primary objective in most social networks, we assume that nodes do not alter their edges with other real nodes as part of a false-name manipulation due to lack of sufficient incentive. That said, alterations to edges with real nodes are a powerful form of manipulation, and preventing such manipulations is an interesting theoretical challenge (see Section 6).

Definition 2 (False-Name Manipulations)**.** *In an instance (G, S, v^*), a voter $v \in S$ can perform a false-name manipulation by creating a set of false nodes M, and edges between a subset of pairs of nodes in $M' \times M'$, where $M' = M \cup \{v\}$. Also, v can choose a subset of nodes in M' to act as voters, and can choose their recommendations. The resulting instance is given by (G', S', v^*), where $V(G') = V \cup M$, $S' \cap (V \setminus \{v\}) = S \setminus \{v\}$, and $E' \cap (V \times V) = E$.*

A node that has a personal opinion may choose to abstain from voting as part of a manipulation. Such a node would be a voter in the underlying true instance, but not in the manipulated instance observed by the mechanism. We refer to it as an "opinionated node" to avoid confusing it with "voters" in the observed instance. In this paper, we only focus on false-name manipulations by individual nodes; Section 6 briefly discusses group false-name manipulations.

Optimal Weight Vector: Preventing false-name manipulations may prohibit us from always choosing the most desirable weight vector for the setting at hand. For the purpose of this paper, we assume a setting in which the ideal weight vector has equal weights for all voters, i.e., in the ideal weight vector each voter in S has weight $1/|S|$. This is interesting due to multiple reasons:

- False-name-proof recommendation mechanisms can employ the knowledge of the social network structure in two clearly distinct ways: i) to weight nodes in a way that provides no incentive for false-name manipulations, and ii) to weight nodes to reflect their level of homophily[2] or trust with the target node. Treating the uniform weight vector as the ideal focuses exclusively on the former purpose. This is also the appropriate choice for networks where no prior information about user opinions is available to conclude homophily.
- Our model also applies to the case where the goal is not to find a recommendation for the target node; instead, the target node conducts a vote on the network, and invites its peers to vote. In this case, treating all voters equally is the de facto fairness consideration in the voting literature.
- Finally, if the opinions offered by the individuals are not subjective preferences, but rather independent noisy estimates of an objective ground truth, using equal weights to aggregate these independent estimates provably yields the most accurate estimation of the underlying ground truth.

That said, aggregating subjective opinions of nodes into a recommendation by weighting the nodes according to their homophily (closeness of opinion) with the target node is an interesting and widely studied topic. As we discuss in Section 6, our results have interesting implications about designing false-name-proof recommendation mechanisms when a model of homophily is given; in this sense, we also view our paper as a stepping stone for studying this more general setting.

Finally, we impose a mild restriction — *symmetry* — on the weight-selecting mechanism. Informally, this requires the mechanism to assign equal weight at least to the nodes that are "symmetrically placed" in the network with respect to the target node.

Definition 3 (Symmetric Mechanisms)**.** *We call a weight-selecting mechanism* symmetric *if, given an instance (G, S, v^*), it assigns equal weights to voters v_1 and v_2 whenever there exists an automorphism of G (i.e., an isomorphism from G to itself) that fixes v^* and maps v_1 to v_2.*

Unless stated otherwise, throughout the paper we will assume a weight-selecting mechanism to be symmetric.

3 Uniform Aggregation

In the context of this paper, the ideal weight-selecting mechanism returns the weight vector that has equal weight for all voters. However, this mechanism suffers from a crucial problem. A node that performs a false-name manipulation by creating an arbitrarily large number of fake nodes and voting through them can accrue a weight arbitrarily close to 1, thus becoming a dictator. Crucially, this manipulation also hurts the other nodes in the network by reducing their weights. To design a robust mechanism, we require that this should not be possible.

Definition 4 (No Harm Axiom)**.** *We say that a weight-selecting mechanism satisfies the no harm axiom if a false-name manipulation by a node does not reduce the weight of any other node in the network. Let $\mathcal{M}^{NH}$ denote the family of symmetric weight-selecting mechanisms satisfying the no harm axiom.*

3.1 No Harm Versus False-Name-Proofness

The standard desideratum in the literature on false-name manipulations is *false-name-proofness*, which requires that even with full information an agent should not be able to find a beneficial false-name manipulation. In our setting, this means a voter should not be able to move the recommendation closer to its personal opinion through a false-name ma-

[2]Homophily is a commonly observed phenomenon where nodes closer in a network are more likely to agree on opinions.

nipulation even if the voter knows the network G, the set of voters S, their personal opinions $\mathbf{r}$, and the target node v^*.

The no harm axiom directly implies that a voter cannot gain weight by performing a false-name manipulation. Could the voter, however, increase the weights of other voters with similar opinions, thereby achieving a more favorable recommendation? This of course depends on how the recommendation mechanism uses the weights to aggregate the opinions. We show that for the weighted median aggregation, the no harm axiom implies false-name-proofness.

Theorem 1. *With real-valued opinions and aggregate recommendation, computing the weighted median of the opinions using the weights returned by a weight-selecting mechanism satisfying the no harm axiom is false-name-proof.*

Proof. Let (G, S, v^*) be the true instance for which the weight vector is $\mathbf{w}$ and the recommendation is x. Suppose $v \in S$ performs a false-name manipulation, after which the weight vector becomes $\mathbf{w}'$ and the recommendation becomes x'. If $r_v = x$, then v has nothing to gain. Without loss of generality, let $r_v > x$. Define $T = \{u \in S | r_u \leq x\}$. Let $w(T) = \sum_{u \in T} w_u$ and $w'(T) = \sum_{u \in T} w'_u$. Then, by the no harm axiom and the definition of weighted median, we have $w'(T) \geq w(T) \geq 0.5$, which implies $x' \leq x$. Hence, the manipulation is not beneficial to v. ■

Theorem 1 shows that the no harm axiom easily yields false-name-proofness. But at first glance, it may seem too strong if the ultimate goal is false-name-proofness. The following result shows that in a simple setting with binary $(0/1)$ opinions and reasonable weighted aggregation functions (e.g., the weighted average), the no harm axiom is *equivalent* to false-name-proofness.

Theorem 2. *Let the opinions be binary (i.e., in $\{0, 1\}$), and the recommendation be computed using a weighted aggregation function that is strictly monotonically increasing in the total weight of all voters with opinion 1, where the weights are computed using a weight-selecting mechanism M. Then, the recommendation system is false-name-proof if and only if M satisfies the no harm axiom.*

Proof. Suppose M satisfies the no harm axiom. Without loss of generality, consider a voter with opinion 0. If the voter performs a false-name manipulation, none of the real voters with opinion 1 lose weight due to the no harm axiom. Hence, the total weight of all voters with opinion 1 does not decrease after the manipulation. Hence, due to monotonicity of the weighted aggregation function, the recommendation cannot decrease due to the manipulation. That is, no false-name manipulation can be beneficial, implying that the recommendation system is false-name-proof.

Now, suppose that M does not satisfy the no harm axiom. Then, there exists an instance (G, S, v^*), a false-name manipulation by $v \in S$ that results in an instance (G', S', v^*), and a voter $u \in S \setminus \{v\}$ such that under M, voter u receives less weight in (G', S', v^*) than in (G, S, v^*). Suppose in (G, S, v^*) all voters in $S \setminus \{u\}$ vote for 1, and only u votes for 0. Since the weights sum to 1 and u loses weight after the manipulation, the total weight of voters with opinion 1 increases after the manipulation. Strict monotonicity of the weighted aggregation function implies that the manipulation would bring the recommendation closer to v's true opinion, 1. Hence, the recommendation system is not false-name-proof in this case. ■

3.2 Search for a Robust Mechanism

Our starting point is a compelling mechanism proposed by Andersen et al. (Andersen et al., 2008) for binary $(0/1)$ recommendations. Imagine doing a random walk on the social network graph starting from the node v^*,[3] and terminating the walk as soon as a voter is encountered. Then, their mechanism recommends an opinion such that the walk is more likely to terminate on a node having that opinion than terminating on a node having the alternative opinion. We observe that this mechanism, which we denote RANDOMWALK, can be viewed as a weight-selecting mechanism.

Definition 5 (RANDOMWALK)**.** *Given an instance (G, S, v^*), the weight-selecting mechanism* RANDOMWALK *outputs the weight vector $\mathbf{w}$ such that for $v \in S$, w_v is the probability that a random walk starting from v^* encounters v before any other voter.*

Our assumption of G being connected and having at least one voter implies that the weights assigned by RANDOMWALK sum to 1. Also, we assume that the edges of the undirected graph G are essentially bidirectional, that is, a walk can traverse an edge in either direction. Crucially, observe that RANDOMWALK satisfies the no harm axiom: Fix a voter v and a walk that leads the random walk to v. When a voter $v' \neq v$ performs a false-name manipulation, the neighborhoods of nodes on the walk do not change. Hence, the walk still leads the random walk to v with the same probability post-manipulation. As this argument applies to every walk leading the random walk to v in the original graph, the total probability of the random walk terminating on v does not reduce after the manipulation. It is also clear that RANDOMWALK is symmetric.

Theorem 3. RANDOMWALK *is a symmetric weight-selecting mechanism satisfying the no harm axiom.*

Example 1. Let G be the network shown in Figure 1. Here, filled nodes represent voters. It is evident that neither v_1 nor v_2 is fake (i.e., they cannot be artificial nodes created by a single node in the network through a false-name manipulation). Recall that by our assumption, v^* does not manipulate. Hence, for uniform aggregation we should weight them equally, if possible.

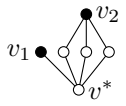

Fig. 1: Graph G

Under RANDOMWALK, voters v_1 and v_2 receive (unequal) weights $2/5$ and $3/5$, respectively. This can be shown by solving systems of linear equations (see Section 4). Note that these probabilities are not $1/4$ and $3/4$, respectively, due

[3] That is, in each step, move from the current vertex to one of its neighbors chosen uniformly at random.

to walks that go from v^* to one of its three non-voter neighbors and return to v^* a number of times, before finally going to v_1.

Admittedly, Andersen et al. (Andersen et al., 2008) study a slightly different setting than ours. Their ultimate goal, unlike ours, is not to uniformly aggregate the opinions; they want the opinion of a voter to be weighted by the level of "trust" v^* can plausibly have for the voter. Hence, in their setting it makes sense to weight the two voters unequally. In other words, our goal is *not* to evaluate RANDOMWALK in our setting, because RANDOMWALK is not designed to give equal weight to voters in the first place. We use Example 1 simply to demonstrate the need to investigate whether there exists a mechanism satisfying the no harm axiom that can provide more uniform weights.

3.3 An Impossibility Result

As the no harm axiom prohibits always selecting the uniform weight vector (with equal weight for all voters), our goal is to find a weight vector that is *as uniform as possible*. To formalize the notion of "uniformity", we use the classic leximin criterion that compares two weight vectors by their minimum weights (and prefers the one with greater minimum weight), and then breaks ties by comparing their second minimum weights, and so on. For example, according to the leximin criterion weight vector $(0.3, 0.5, 0.2)$ is better (i.e., more uniform) than weight vector $(0.4, 0.5, 0.1)$, but is no different than weight vector $(0.5, 0.3, 0.2)$. The leximin criterion has been studied extensively in the literature (Sen, 1970; Moulin, 1991, 2003), and has been applied successfully in a broad spectrum of domains including constraint programming (Bouveret and Lemaître, 2009), wireless networks (Huang and Bensaou, 2001), resource allocation (Ghodsi et al., 2011; Kurokawa, Procaccia, and Shah, 2015), cake-cutting (Chen et al., 2013), and kidney exchange (Roth, Sönmez, and Ünver, 2005).

Definition 6 (Leximin Comparison)**.** *On an instance (G, S, v^*), let weight-selecting mechanisms M and M' return weight vectors $\mathbf{w}$ and $\mathbf{w}'$, consisting of weights $(w_1, \ldots, w_{|S|})$ and $(w'_1, \ldots, w'_{|S|})$, respectively, sorted in the non-decreasing order. Then, M is* leximin-better *than M' on (G, S, v^*) if there exists $t \in \{1, \ldots, |S|\}$ such that $w_i = w'_i$ for all $i \in \{1, \ldots, t-1\}$ and $w_t > w'_t$.*

Comparing mechanisms across instances, we say that M is leximin-better than M' if M' is not leximin-better than M on any instance, and M is leximin-better than M' on at least one instance.

Definition 7 (Leximin-Optimality)**.** *In a family of weight-selecting mechanisms $\mathcal{C}$, mechanism $M \in \mathcal{C}$ is called leximin-optimal for $\mathcal{C}$ if M is leximin-better than every other mechanism in $\mathcal{C}$.*

We can now cast our search for a good mechanism as a formal question. *Does there exist a mechanism that is leximin-optimal for the family $\mathcal{M}^{NH}$ of symmetric weight-selecting mechanisms satisfying the no harm axiom?* Note that at most one mechanism could satisfy this desideratum. Unfortunately, the next result shows that in our case none meets the bar.

Theorem 4. *No mechanism is leximin-optimal for $\mathcal{M}^{NH}$.*

Proof. Let G_1 and G_2 be the networks shown in Figure 2. Suppose for contradiction that there exists a weight-

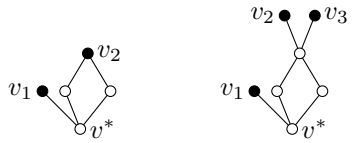

(a) Graph G_1 (b) Graph G_2

Fig. 2: Impossibility of leximin-optimality for $\mathcal{M}^{\text{NH}}$

selecting mechanism $M \in \mathcal{M}^{\text{NH}}$ that is leximin-optimal for $\mathcal{M}^{\text{NH}}$. It can be shown that there exists a mechanism in $\mathcal{M}^{\text{NH}}$ that weights both voters in G_1 equally. While RANDOMWALK does not satisfy this, the reader may check that the mechanism LEGIT$^+$ that we later propose in Section 3.4 does. Leximin-optimality of M now implies that M must assign weight $1/2$ to both voters in G_1.

Next, note that G_2 is created when voter v_2 in G_1 performs a false-name manipulation. Thus, the no harm axiom implies that M must still assign a weight of at least $1/2$ to v_1 in G_2. As the remaining weight is divided equally among the remaining two voters by symmetry, it follows that the minimum weight in G_2 under M is at most $1/4$. However, it can be checked that the minimum weight in G_2 under RANDOMWALK is $2/5$, which is greater than $1/4$. Hence, RANDOMWALK $\in \mathcal{M}^{\text{NH}}$ is leximin-better than M on G_2, which contradicts leximin-optimality of M for $\mathcal{M}^{\text{NH}}$. ■

3.4 A Possibility Result

Theorem 4 implies that subject to the no harm axiom, a mechanism cannot be the best on every instance. It faces an inevitable trade-off whereby choosing to be better on one instance requires it to be worse on another. Which instances should get more emphasis? In social networks, often very few users make the effort to vote, and this scarcity of information is further exacerbated in smaller networks. Thus, arguably, achieving a uniform weight vector is more important in smaller networks so that every opinion counts. In larger networks, it is often excusable to discard a few opinions in order to achieve robustness. We translate this informal goal of giving more importance to smaller networks into a formal desideratum, *"optimized for small networks"*, which we view as a novel conceptual contribution of the paper as its formulation may be useful in other settings as well.

Definition 8 (Domination)**.** *For weight-selecting mechanisms M and M', we say that M* dominates *M' on network G for target node v^* if M' is not leximin-better than M on (G, S, v^*) for any $S \subseteq V(G) \backslash \{v^*\}$, and M is leximin-better than M' on (G, S, v^*) for some $S \subseteq V(G) \setminus \{v^*\}$.*

Definition 9 (Optimized for Small Networks)**.** *For a family of weight-selecting mechanisms $\mathcal{C}$ and $M \in \mathcal{C}$, we say that M is optimized for small networks within $\mathcal{C}$ if the following holds: If $M' \in \mathcal{C}$ is leximin-better than M on an instance*

(G, S, v^*), *there exists a strict subgraph* H *of* G *with* $v^* \in V(H)$ *such that* M *dominates* M' *on* H *for target node* v^*.

While being optimized for small networks is weaker than (is implied by) leximin-optimality, the bar is still high, as the next observation shows. Its simple proof is omitted due to space constraints.

Proposition 1. *In a family of weight-selecting mechanisms, at most one mechanism is optimized for small networks.*

We now design an intuitive weight-selecting mechanism, and show that it is optimized for small networks within the family $\mathcal{M}^{\text{NH}}$. A key idea behind the mechanism is due to Conitzer et al. (Conitzer et al., 2010), who propose a method of identifying *certifiably legitimate nodes* in a network in the presence of false-name manipulations. In their more general setting, this is a tricky problem, but in our setting it boils down to a simple observation: v *could possibly be a fake node created by* u *if and only if removing* u *disconnects* v *from* v^*.

Let $F(u)$, called the *lobe* of u, be the set of all nodes that become disconnected from v^* by removing u. As a convention, $u \notin F(u)$, and $F(v^*)$ is undefined. Now, node v is certifiably legitimate if $v \notin F(u)$ for any $u \in V \setminus \{v^*\}$. In other words, v should remain connected to v^* after removing any single node. Equivalently, it should either be a direct neighbor of v^*, or be 2-vertex-connected to v^*.[4] Suppose we can weight all certifiably legitimate voters equally. The following lemma helps us deal with the remaining voters.

Lemma 1. *A symmetric weight-selecting mechanism satisfying the no harm axiom cannot assign a positive weight to any node in the lobe of a voter.*

Proof. Suppose for contradiction that node $v \in F(u)$ receives weight $\delta > 0$, where u is a voter. Suppose all nodes in the graph (including those in $F(u)$) are real. As the nodes in $F(u)$ are only connected to the remaining network through u, under a false-name manipulation u can create N copies of $F(u)$ that are attached to u in a way identical to how $F(u)$ is attached. The no harm axiom implies that v still receives weight at least δ, and by symmetry, now so does each of its N copies. However, this is infeasible when $N > 1/\delta$ as the weights must sum to 1. ■

While Lemma 1 requires us to discard all nodes in the lobe of a voter, it does not prevent us from distributing the weight that a certifiably legitimate non-voter would have received (had it been a voter) to the nodes in its lobe. In fact, such distribution is necessary for the weights to sum to 1 when all voters reside in lobes of other nodes. A natural way is to apply this procedure recursively in each such lobe. This leads to our mechanism, which we call LEGIT$^+$ because it recursively passes legitimacy to nodes. It is presented as Algorithm 1. Crucially, we only recursively apply the mechanism to a lobe if it has a voter. Note that our mechanism assigns a positive weight to the maximal set of nodes subject to Lemma 1. This leads us to the next result.

[4]Using Menger's theorem (Menger, 1927), being 2-vertex-connected to v^* is equivalent to having two vertex-disjoint paths to v^*.

ALGORITHM 1: Mechanism LEGIT$^+$

Data: Social network G, set of voters S, central node v^*
Result: Weight vector $\mathbf{w} = (w_v)_{v \in S}$
$\forall u \in V \setminus \{v^*\}, F(u) \leftarrow \big\{t \in V \setminus \{u\} \;\big|\; t$ is not connected to v^* in $G_{V \setminus \{u\}}\big\}$;
$L \leftarrow \big\{v \in V \setminus \{v^*\} \big|\; (\nexists u \in V \setminus \{v^*\} : v \in F(u)) \wedge (F(v) \cup \{v\}) \cap S \neq \emptyset\big\}$;
$\forall v \in S, w_v \leftarrow 0$;
for $v \in L$ **do**
  if $v \in S$ **then**
    $w_v \leftarrow 1/|L|$;
  else
    $T \leftarrow F(v) \cup \{v\}$;
    $\mathbf{w}^{\mathbf{rec}} \leftarrow$ LEGIT$^+(G_T, S \cap T, v)$;
    for $u \in F(v) \cap S$ **do** $w_u \leftarrow w_u^{\text{rec}} \cdot 1/|L|$;
  end
end
return $\mathbf{w} = (w_v)_{v \in S}$;

Lemma 2. *If a node receives zero weight under* LEGIT$^+$*, it receives zero weight under all mechanisms in* $\mathcal{M}^{NH}$.

We are now ready for the main result of this paper.

Theorem 5. LEGIT$^+$ *is optimized for small networks within the family* $\mathcal{M}^{NH}$.

4 Computational Complexity

Let us begin with RANDOMWALK. Andersen et al. (2008) show that aggregation of binary recommendations under RANDOMWALK amounts to solving a single system of linear equations $Ax = b$, where the LHS matrix A is $n \times n$ ($n = |V|$ is the number of nodes) and the RHS vector b is $n \times 1$. Solving this system can take, even with recent exact solvers, $O(|V|^{1.5} \cdot (|V| + |E|))$ time (Eberly et al., 2006). Implementing RANDOMWALK as a weight-selecting mechanism is computationally even more difficult. We need to solve one system of linear equations for each voter, which can take $O(|V|^{2.5} \cdot (|V| + |E|))$ time (Eberly et al., 2006).

Let us now consider LEGIT$^+$. Arguably, it is harder to describe than RANDOMWALK, and Algorithm 1 is more intricate than simply solving a collection of linear systems. More specifically, in the first step of Algorithm 1 simply computing $F(u)$ for every node u would naïvely take $O(|V| \cdot (|V| + |E|))$ time. Surprisingly, we show that there exists a more efficient implementation that computes the weights under LEGIT$^+$ in merely $O(|V|+|E|)$ (linear) time. This implementation uses as a subroutine the remarkable linear time algorithm by Hopcroft and Tarjan (1973) for finding biconnected components in a graph. A *biconnected component* (or a *block*) is a maximal 2-vertex-connected subgraph. Nodes that belong to multiple blocks (i.e., whose removal disconnects the graph) are called *cut vertices* or *articulation points*. A connected graph G decomposes into a *block-cut tree* T whose vertices are the blocks and the articulation points of G, and a block B and an articulation point u are connected if $u \in B$.

Let A denote the set of articulation points of G, and $\mathcal{B}_u$ denote the set of blocks of G containing u. First, u has a non-empty lobe $F(u)$ if and only if $u \in A$. Next, if $u \in A$, the lobe $F(u)$ can be computed as follows. Remove the vertex of T representing u, which disconnects T into connected components, one of which contains all blocks containing v^*. The set of nodes in the blocks contained in every other connected component of the tree (except u itself) constitute $F(u)$. This key observation leads us to a linear time implementation of LEGIT$^+$, the details of which are omitted due to space constraints.

Theorem 6. *Weights under* LEGIT$^+$ *can be computed in* $O(|V| + |E|)$ *time.*

5 Experiments

We compare LEGIT$^+$ with two baseline mechanisms: LEGIT and RANDOMWALK. We define weight-selecting mechanism LEGIT as the simpler version of LEGIT$^+$ that assigns equal weight to all certifiably legitimate voters, but does *not* apply the procedure recursively within the lobes of certifiably legitimate non-voters. Thus, comparison with LEGIT indicates the gain from recursively applying LEGIT$^+$ within the lobes of certifiably legitimate non-voters. We note that LEGIT$^+$ is expected to (though theoretically not guaranteed to) outperform RANDOMWALK, because RANDOMWALK is not designed to assign uniform weights.

We perform experiments using 16 real-world social networks from the KONECT project (Kunegis, 2013). The number of nodes and edges in these networks vary from 23 to 26,475, and from 78 to 146,385, respectively.[5] For each network G, we sample the target node v^* uniformly at random. For each pair (G, v^*), we determine the set of voters by making each node in the network a voter independently with probability p_{vote}. We use both low values (from 0.01 to 0.09 in increments of 0.02) and high values (from 0.1 to 0.9 in increments of 0.2) of p_{vote}, representative of varying levels of voter engagement. For each network and each of 10 values of p_{vote}, we choose 100 random target nodes, and for each target node, choose 100 random subsets of voters. In the results presented below, we compare LEGIT$^+$ with LEGIT and RANDOMWALK across the simulations for each network. To solve the linear system in RANDOMWALK, we use Matlab's mldivide operator, and to find the biconnected components in LEGIT$^+$, we use the MatlabBGL library [6].

Figure 3(a) shows a log-log plot of the running time of all three mechanisms (LEGIT$^+$ as magenta diamonds, LEGIT as red circles, and RANDOMWALK as blue stars) as a function of the number of nodes in the network. The experiments were performed on a dual-core machine with 3.10 GHz processors and 8 GB RAM. While LEGIT is trivially faster than LEGIT$^+$ (as it requires a strictly less number of operations), the difference is not significant. On the other hand, while RANDOMWALK is slightly faster than LEGIT$^+$ on smaller networks, LEGIT$^+$ is significantly faster on networks with more than 200 nodes. This is consistent with our result from Section 4 that the worst-case complexity is significantly lower for LEGIT$^+$ than for RANDOMWALK (linear versus super-quadratic). Across the entire experiment, LEGIT$^+$ ran about 13 times faster than RANDOMWALK, and only about 3 times slower than LEGIT.

In the remaining figures, we only plot two lines: one that compares LEGIT$^+$ with LEGIT (with red circles), and one that compares LEGIT$^+$ with RANDOMWALK (with blue stars).

Our next goal is to determine which mechanism outputs a more uniform weight vector. Lacking an objective definition of uniformity, we use three metrics: i) leximin comparison as used in our theoretical results in Section 3, ii) the percentage of voters discarded, i.e., assigned zero weight to (the lower, the better), iii) the (L^2-)distance from the uniform weight vector, which is equal to the variance of the weight vector (the lower, the better).

Figure 3(b) shows that LEGIT$^+$ is leximin-better than both LEGIT and RANDOMWALK in more than 50% simulations in each network. In fact, it is leximin-better than LEGIT (resp. RANDOMWALK) in more than 75% (resp. 85%) simulations in all but one (resp. two) networks. Superior empirical performance in such large networks nicely complements our theoretical result (Theorem 5), which indicates that LEGIT$^+$ should be superior in small networks in general.

Next, while Lemma 2 ensures that LEGIT$^+$ discards the smallest subset of voters subject to the no harm axiom, Figure 3(c) shows that LEGIT and RANDOMWALK can discard up to 60% and 20% more voters, respectively, than LEGIT$^+$ (about 30% and 10%, respectively, on average across networks).

Finally, comparing variance of the returned weight vector, Figure 3(d) shows that LEGIT$^+$ performs better than both LEGIT and RANDOMWALK in more than 50% simulations in each network. Further, it outperforms LEGIT in at least 69% simulations in all but one network, and RANDOMWALK in at least 89% simulations in all but one network.

So far we have focused on the setting where the opinions of voters are subjective, and the goal is to find a weight vector as close to uniform as possible. We now present empirical results for a slightly different setting in which there exists a binary (0/1) ground truth, and the goal is to pinpoint it by aggregating binary opinions of voters, each of which is "correct" with probability $p_{\text{acc}} > 0.5$. The accuracy of a weight-selecting mechanism on an instance (G, S, v^*) is the probability that the mechanism assigns higher total weight to voters with the correct opinion than to voters with the incorrect opinion. While we do not have theoretical results for this setting, we can evaluate the

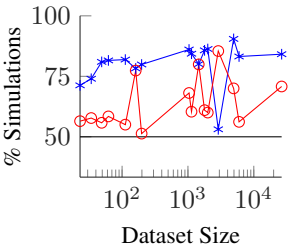

Fig. 4: Accuracy

[5] Running experiments on the larger datasets was infeasible due to the prohibitive running time of RANDOMWALK.

[6] https://www.cs.purdue.edu/homes/dgleich/packages/matlab_bgl/

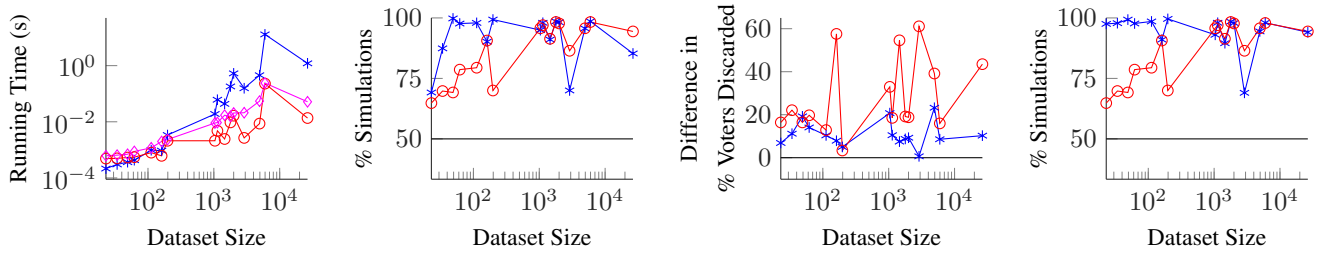

(a) Running Time (b) Leximin Comparison (c) Diff. in % Voters Discarded (d) Comparing Variance

Fig. 3: Comparison of LEGIT$^+$ with RANDOMWALK on real-world social networks

mechanisms empirically. For p_{acc}, we use both low values (0.51 to 0.59 in increments of 0.02) and high values (0.6 to 0.9 in increments of 0.1).

Figure 4 shows that LEGIT$^+$ achieves better accuracy than both LEGIT and RANDOMWALK in more than 50% simulations in each network. Also, note that LEGIT$^+$ achieves better accuracy than RANDOMWALK in at least 70% simulations in all but one network.

6 Discussion

Recall the median-of-medians rule from the introduction: the recommendation is the median of the opinions of the target agent's friends, and for a friend who does not provide an opinion, we construct one by taking the median of *his* friends' opinions, and so on. In conjunction with the weighted median aggregation rule (as in Theorem 1), LEGIT$^+$ can be seen as a similar rule, "median-of-medians for legitimate nodes": instead of taking the median of friends' opinions, take the median of the opinions of (certifiably) legitimate nodes, and for such nodes that do not provide an opinion, construct one recursively from opinions in their lobes.

We uniquely characterize LEGIT$^+$ within the family of symmetric weight-selecting mechanisms satisfying the no harm axiom. We show this axiom to be closely related, but in the most general setting incomparable, to false-name-proofness. The no harm axiom is only defined for weight-selecting mechanisms. It remains to be seen whether we can pinpoint an overall recommendation mechanism that uses LEGIT$^+$ (e.g., median-of-median for legitimate nodes) within the more general family of false-name-proof mechanisms.

Importantly, in this paper we consider the uniform weight vector as idealistic. In the context of aggregating subjective opinions into a personal recommendation for the target node, this only makes sense in the absence of knowledge of correlation among user preferences (e.g., homophily of opinions). However, note that Lemma 1 provides a necessary condition for satisfying the no harm axiom — in the form of having to assign zero weight to specific nodes — *even in the presence of homophily*. Given a model of homophily, we must start by assigning zero weight to such nodes. Weighting the remaining nodes to maximally align the recommendation with the target node's preference is still a difficult problem. Fortunately, it can be shown that choosing the remaining weights as a function of the vertex-connectivity of a node to the target node is sufficient to guarantee the no harm axiom. However, this approach is likely to be suboptimal. An immediate next step is to design better ways of incorporating homophily subject to the no harm axiom.

An interesting direction for future research is to study stronger manipulations. For example, LEGIT$^+$ does not prevent group false-name manipulations or manipulations where nodes may delete their existing edges with other real nodes. Can we effectively prevent them? While RANDOMWALK is group false-name-proof, it can be shown that it is not optimized for small networks among symmetric group false-name-proof mechanisms. Does there exist such a mechanism (recall that there can be at most one)?

False-name manipulations are an increasingly serious concern in social networks, especially with the effortless accessibility and increasing popularity of automated tools for creating fake accounts (Pathak, 2014). Given the difficulty of distinguishing fake accounts from real ones, we believe that the study of false-name-proofness is the key to building the next generation of reliable recommendation systems.

Acknowledgements

We are thankful for support from NSF under awards IIS-1527434, IIS-0953756, CCF-1101659, and CCF-1337215, ARO under grants W911NF-12-1-0550 and W911NF-11-1-0332, ERC under StG 639945 (ACCORD), a Guggenheim Fellowship, and a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation. This work was done in part while Conitzer was visiting the Simons Institute for the Theory of Computing.

References

Adomavicius, G., and Tuzhilin, A. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17(6):734–749.

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 *Proceedings of the 17th International World Wide Web Conference (WWW)*, 199–208.

Balabanović, M., and Shoham, Y. 1997. Fab: content-based, collaborative recommendation. *Communications of the ACM* 40(3):66–72.

Barbera, S.; Sonnenschein, H.; and Zhou, L. 1991. Voting by committees. *Econometrica* 595–609.

Bennett, J., and Lanning, S. 2007. The Netflix prize. In *Proceedings of KDD cup and workshop*, 35.

Border, K., and Jordan, J. 1983. Straightforward elections, unanimity and phantom voters. *Review of Economic Studies* 50:153–170.

Bouveret, S., and Lemaître, M. 2009. Computing leximin-optimal solutions in constraint networks. *Artificial Intelligence* 173(2):343–364.

Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior* 77:284–297.

Conitzer, V.; Immorlica, N.; Letchford, J.; Munagala, K.; and Wagman, L. 2010. False-name-proofness in social networks. In *Proceedings of the 6th International Workshop on Internet and Network Economics (WINE)*, 209–221.

Conitzer, V. 2008. Anonymity-proof voting rules. In *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE)*, 295–306.

Eberly, W.; Giesbrecht, M.; Giorgi, P.; Storjohann, A.; and Villard, G. 2006. Solving sparse rational linear systems. In *Proceedings of the 2006 International Symposium on Symbolic and Algebraic Computation (ISSAC)*, 63–70. ACM.

Ghodsi, A.; Zaharia, M.; Hindman, B.; Konwinski, A.; Shenker, S.; and Stoica, I. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, 24–37.

Golbeck, J. 2006. *Generating predictive movie recommendations from trust in social networks*. Springer.

He, J., and Chu, W. W. 2010. *A social network-based recommender system (SNRS)*. Springer.

Hopcroft, J., and Tarjan, R. 1973. Algorithm 447: Efficient algorithms for graph manipulation. *Communications of the ACM* 16(6):372–378.

Huang, X. L., and Bensaou, B. 2001. On max-min fairness and scheduling in wireless ad-hoc networks: analytical framework and implementation. In *Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc)*, 221–231.

Iwasaki, A.; Conitzer, V.; Omori, Y.; Sakurai, Y.; Todo, T.; Guo, M.; and Yokoo, M. 2010. Worst-case efficiency ratio in false-name-proof combinatorial auction mechanisms. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 633–640.

Kunegis, J. 2013. KONECT - The Koblenz network collection. In *Proceedings of the Int. Web Observatory Workshop*.

Kurokawa, D.; Procaccia, A. D.; and Shah, N. 2015. Leximin allocations in the real world. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, 345–362.

Menger, K. 1927. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae* 1(10):96–115.

Moulin, H. 1980. On strategy-proofness and single-peakedness. *Public Choice* 35:437–455.

Moulin, H. 1991. *Axioms of cooperative decision making*. Cambridge University Press.

Moulin, H. 2003. *Fair Division and Collective Welfare*. MIT Press.

Pathak, A. 2014. *An analysis of various tools, methods and systems to generate fake accounts for social media*. Ph.D. Dissertation, Northeastern University Boston.

Pazzani, M., and Billsus, D. 2007. Content-based recommendation systems. *The Adaptive Web* 325–341.

Ricci, F.; Rokach, L.; and Shapira, B. 2011. *Introduction to recommender systems handbook*. Springer.

Roth, A. E.; Sönmez, T.; and Ünver, M. U. 2005. Pairwise kidney exchange. *Journal of Economic Theory* 125:151–188.

Sen, A. 1970. *Collective Choice and Social Welfare*. North-Holland.

Todo, T., and Conitzer, V. 2013. False-name-proof matching. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 311–318.

Todo, T.; Iwasaki, A.; Yokoo, M.; and Sakurai, Y. 2009. Characterizing false-name-proof allocation rules in combinatorial auctions. In *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 265–272.

Todo, T.; Iwasaki, A.; and Yokoo, M. 2010. False-name-proofness in facility location problem on the real line. In *Proceedings of the 6th International Workshop on Internet and Network Economics (WINE)*, 559–562.

Todo, T.; Iwasaki, A.; and Yokoo, M. 2011. False-name-proof mechanism design without money. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 651–658.

Wagman, L., and Conitzer, V. 2008. Optimal false-name-proof voting rules with costly voting. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI)*, 190–195.

Yokoo, M.; Sakurai, Y.; and Matsubara, S. 2001. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence* 130(2):167–181.

Yokoo, M.; Sakurai, Y.; and Matsubara, S. 2004. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. *Games and Economic Behavior* 46(1):174–188.

Yokoo, M. 2003. Characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, 733–739.