

Sum of Us: Strategyproof Selection from the Selectors

Noga Alon
School of Mathematical Sciences
Tel Aviv University and
Microsoft Israel R&D Center
nogaa@tau.ac.il

Felix Fischer
School of Engineering and
Applied Sciences
Harvard University
fischerf@seas.harvard.edu

Ariel Procaccia
School of Engineering and
Applied Sciences
Harvard University
arielpro@seas.harvard.edu

Moshe Tennenholtz
Microsoft Israel R&D Center and
Faculty of Industrial Engineering and Management
Technion – Israel Institute of Technology
moshet@microsoft.com

ABSTRACT

We consider the special case of approval voting when the set of agents and the set of alternatives coincide. This captures situations in which the members of an organization want to elect a president or a committee from their ranks, as well as a variety of problems in networked environments, for example in internet search, social networks like Twitter, or reputation systems like Epinions. More precisely, we look at a setting where each member of a set of n agents approves or disapproves of any other member of the set and we want to select a subset of k agents, for a given value of k , in a strategyproof and approximately efficient way. Here, strategyproofness means that no agent can improve its own chances of being selected by changing the set of other agents it approves. A mechanism is said to provide an approximation ratio of α for some $\alpha \geq 1$ if the ratio between the sum of approval scores of any set of size k and that of the set selected by the mechanism is always at most α . We show that for $k \in \{1, 2, \dots, n-1\}$, no deterministic strategyproof mechanism can provide a finite approximation ratio. We then present a randomized strategyproof mechanism that provides an approximation ratio that is bounded from above by four for any value of k , and approaches one as k grows.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Economics, Theory

Keywords

Social choice, Approval voting, Approximate mechanism design without money

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

TARK XIII, Jul 12–14, Groningen, The Netherlands

Copyright 2011 ACM 978-1-4503-0707-9/11/07 ...\$10.00.

1. INTRODUCTION

One of the most well-studied settings in social choice theory concerns a set of *agents* (also known as *voters* or *individuals*) and a set of *alternatives* (also known as *candidates*). The agents express their preferences over the alternatives, and these are mapped by some function to a winning alternative or set of winning alternatives. In one prominent variation, each agent must select a subset of alternatives it approves; this setting is known as *approval voting* [5].

We consider the special case of approval voting when the set of agents and the set of alternatives coincide; this for example occurs when the members of an organization use approval voting to elect a president or a committee from their ranks.¹ We model this situation by a directed graph on the set of agents, where an edge from agent i to agent j means that agent i approves, votes for, trusts, or supports agent j . Our goal is to select a subset of k “best” agents for each graph and a given value of k , and we will elaborate on what we mean by “best” momentarily.

The fact that agents and alternatives coincide allows us to make additional assumptions about agents’ preferences. Indeed, we will assume that each agent is only interested in whether it is among those selected, i.e., that it receives utility one if selected and zero otherwise. We will see, however, that our results in fact hold for any setting where agents give their own selection priority over that of their approved candidates. This assumption, which is very reasonable in practice, is discussed in more detail in Section 5.

A deterministic *k-selection mechanism* is a function that maps a given graph on the set of agents to a k -subset of selected agents. We also consider *randomized k-selection mechanisms*, which randomly select a subset. The set of other agents approved by a particular agent, i.e., its set of outgoing edges in the graph, is private information of that agent. Agents thus play a game: each of them reports a set of outgoing edges, which might differ from the true one, and the mechanism selects a subset of agents based on the reported edges. We say that a mechanism is *strategyproof*

¹Approval voting is employed in this exact context for example by scientific organizations such as the American Mathematical Society (AMS), the Institute of Electrical and Electronics Engineers (IEEE), the Game Theory Society (GTS), and the International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS).

(*SP*) if an agent cannot increase its chances of being selected by misreporting its outgoing edges, even if it has complete information about the rest of the graph. We further say that a mechanism is *group strategyproof (GSP)* if even a coalition of agents cannot all gain from misreporting their outgoing edges.

It remains to be specified what we mean by selecting the “best” agents. In this paper, we measure the quality of a set of agents by their total number of incoming edges, i.e., the sum of their indegrees. Our goal will be to optimize this target function. Note that this goal is in a sense orthogonal to the agents’ interests, which may make the design of good SP mechanisms difficult.

In addition to traditional voting settings, our model also captures different problems in networked environments. Consider for example an Internet search setting where agents correspond to web sites and edges represent hyperlinks. A search engine should return the top k web sites for a given graph. Each specific web site, or more accurately its webmaster, is naturally concerned with appearing at the top of the search results, and to this end may add or remove hyperlinks at will.

Another example can be found in the context of social networks. While some social networks, like Facebook,² are associated with undirected graphs, there are many examples with unilateral connections. Users of the reputation system Epinions³ unilaterally choose which other users to trust, thus establishing a “Web of Trust.” In the the social network Twitter,⁴ which of late has become wildly popular, a user chooses which other users to “follow.” In these “directed” social networks, choosing a k -subset with maximum overall indegree simply means selecting the k most popular or most trusted users. Applications include setting up a committee, recommending a trusted group of vendors, targeting a group for an advertising campaign, or simply holding a popularity contest. The last point may seem like pure fantasy, but quite recently celebrity users of Twitter in fact held a race to the milestone of one million followers; the dubious honor ultimately went to actor Ashton Kutcher. Clearly Mr. Kutcher could increase the chance of winning by not following any other users.

Since a mechanism that selects an optimal subset (in terms of total indegree) is clearly not SP, we will resort to approximate optimality. More precisely, we seek SP mechanisms that provide a good approximation, in the usual sense, to the total indegree. Crucially, approximation is *not* employed in this context to circumvent computational complexity, as the problem of selecting an optimal subset is obviously tractable, but in order to sufficiently broaden the space of acceptable mechanisms to include SP ones.

Context and related work The work in this paper falls squarely into the realm of *approximate mechanism design without money*, an agenda recently introduced by some of us [20] building on earlier work for example by Dekel et al. [8]. This agenda advocates the design of SP approximation mechanisms *without payments* for structured, and preferably computationally tractable, optimization problems. Indeed, while almost all the work in the field of *algorithmic mechanism design* [19] considers mechanisms that

are allowed to transfer payments to and from the agents, monetary transfers are often infeasible due to ethical or legal considerations, like in voting, or for practical reasons like security and accountability, like in the networked environments discussed above (also see, e.g., the book chapter by Schummer and Vohra [22]). Our notion of a mechanism, sometimes referred to as a *social choice function* in the social choice literature, therefore precludes payments by definition.

Recent work by Holzman and Moulin [16] considers a special case of our setting, in the sense that a single agent must be selected, but asks a question that is fundamentally different: rather than seeking to optimize a target function under the strategyproofness constraint (which they refer to as *impartiality*), Holzman and Moulin look for mechanisms that satisfy intuitive axiomatic properties and identify specific families of deterministic mechanisms that are particularly desirable from the axiomatic point of view. They motivate their work by peer evaluation in communities of experts, which nicely complements the voting and social networking scenarios discussed above.

Altman and Tennenholtz [3] perform an axiomatic study of preference aggregation in settings where the set of agents and the set of alternatives coincide and the desired outcome is a ranking of the agents. The case of our setting where a single agent is to be selected also is a special case of so-called *selection games* [2]. Altman et al. [4] study manipulation in the context of mechanisms that select agents based on a tournament, i.e., a directed graph representing pairwise comparisons between the agents.

LeGrand et al. [17] consider approximations of the (less standard) minimax solution for approval voting, which selects alternatives in a way that minimizes the maximum Hamming distance to the agents’ ballots (viewed as binary vectors). They show that the optimization problem is NP-hard, and provide a trivial 3-approximation algorithm that simply chooses the subset that is closest to the ballot of an arbitrary agent. They further observe that this algorithm is SP when an agent’s (dis)utility is its Hamming distance to the selected subset. Caragiannis et al. [6] continue this line of work by designing better polynomial-time approximation algorithms, slightly improving the approximation ratio achievable via SP algorithms, and providing lower bounds.

Finally, our work is related to the literature on manipulation of reputation systems. Reputation systems are often modeled as weighted directed graphs, and a reputation function is then used that maps a given graph to reputation values for the agents (e.g., [7, 12]). While our positive results can be extended to weighted graphs when the target function is the sum of weights on incoming edges, this would hardly be a meaningful target function. Indeed, in this context the absence of a specific incoming edge (indicating lack of knowledge) would be preferable to an edge with low weight (indicating distrust); see Section 5 for further discussion.

Results and techniques We give rather tight upper and lower bounds on the approximation ratio achievable by k -selection mechanisms in the setting described above; the properties of the mechanisms fall along two orthogonal dimensions, deterministic vs. randomized and SP vs. GSP. Table 1 summarizes our results.

We begin by studying deterministic k -selection mechanisms in Section 3. It is quite easy to see that no deterministic SP 1-selection mechanism can provide a finite approximation ratio. Intuitively, the same should not be true

²<http://facebook.com>

³<http://epinions.com>

⁴<http://twitter.com>

		Deterministic	Randomized
SP	UB	∞	$\min\{4, 1 + \mathcal{O}(1/k^{1/3})\}$
	LB	∞	$1 + \Omega(1/k^2)$
GSP	UB	∞	n/k
	LB	∞	$(n-1)/k$

Table 1: Approximation ratio achievable by k -selection mechanisms for n agents. UB stands for upper bound, LB for lower bound, SP for strategyproof, and GSP for group strategyproof.

for larger values of k . Indeed, in order to guarantee a finite approximation ratio, a mechanism should very simply select a subset of agents with at least one incoming edge, if there is such a set. In the extreme case when $k = n - 1$, we must select all the agents save one, and the question is whether there exists an SP mechanism that never eliminates a unique agent with positive indegree. Our first result gives a surprising negative answer to this question, and in fact holds for every value of k .

THEOREM 3.1. *Let $N = \{1, \dots, n\}$, $n \geq 2$, and $k \in \{1, \dots, n-1\}$. Then there is no deterministic SP k -selection mechanism that provides a finite approximation ratio.*

The proof of this result is concise but rather tricky. It involves two main arguments. We first restrict our attention to a subset of the graphs, namely to stars with all edges directed at a specific agent. An SP mechanism for such graphs can be represented using a function over the boolean $(n-1)$ -cube, which must satisfy certain constraints. We then use a parity argument to show that the constraints lead to a contradiction.

In Section 4 we turn to randomized k -selection mechanisms. We design a randomized mechanism, called *Random m -Partition (m -RP)* and parameterized by m , that works by randomly partitioning the set of agents into m subsets, and then selecting the (roughly) k/m agents with largest indegree from each subset, when only the incoming edges from the other subsets are taken into account. This rather simple technique is reminiscent of work on *random sampling* in the context of auctions for digital goods [11, 15, 10] and combinatorial auctions [9], although our problem is fundamentally different. We obtain the following theorem.

THEOREM 4.1. *Let $N = \{1, \dots, n\}$, $k \in \{1, \dots, n-1\}$. For every value of m , m -RP is SP. Furthermore,*

1. *2-RP provides an approximation ratio of four, and*
2. *for $k \geq 2$, $\left(\left\lceil k^{1/3} \right\rceil\right)$ -RP provides an approximation ratio of $1 + \mathcal{O}(1/k^{1/3})$.*

For a given number k of agents to be selected, we can in fact choose the best value of m when applying m -RP. There thus exists a mechanism with an approximation ratio that is bounded from above by four for any value of k , and approaches one as k grows.

In addition, we prove a lower bound of $1 + \Omega(1/k^2)$ on the approximation ratio achievable by any randomized SP k -selection mechanism; in particular, the lower bound is two for $k = 1$.

As our final result, we obtain a lower bound of $(n-1)/k$ for randomized GSP k -selection mechanisms. This result implies that when asking for group strategyproofness one essentially cannot do better than simply selecting k agents at random, which is obviously GSP and provides an approximation ratio of n/k .

2. THE MODEL

Let $N = \{1, \dots, n\}$ be a set of *agents*. For each $k = 1, \dots, n$, let $\mathcal{S}_k = \mathcal{S}_k(n)$ be the collection of k -subsets of N , i.e., $\mathcal{S}_k = \{S \subseteq N : |S| = k\}$. Agents' preferences are modeled by a directed graph $G = (N, E)$ without self-loops, i.e., $E \subseteq V \times V$ such that for all $(i, j) \in E$, $i \neq j$. The set of such graphs is denoted by $\mathcal{G} = \mathcal{G}(N)$.

A *deterministic k -selection mechanism* is a function $f : \mathcal{G} \rightarrow \mathcal{S}_k$ that for each graph selects a subset of the agents. When the subset $S \subseteq N$ is selected, agent $i \in N$ obtains utility $u_i(S) = 1$ if $i \in S$ and $u_i(S) = 0$ otherwise, i.e., agents only care about whether they are selected or not. We further discuss this utility model in Section 5.

A *randomized k -selection mechanism* is a function $f : \mathcal{G} \rightarrow \Delta(\mathcal{S}_k)$, where $\Delta(\mathcal{S}_k)$ is the set of probability distributions over \mathcal{S}_k . Given a distribution $\mu \in \Delta(\mathcal{S}_k)$, the utility of agent $i \in N$ is

$$u_i(\mu) = \mathbb{E}_{S \sim \mu}[u_i(S)] = \Pr_{S \sim \mu}[i \in S].$$

Deterministic mechanisms are treated as a special case of randomized ones, where for each graph a set of agents is selected with probability one.

We say that a k -selection mechanism is *strategyproof (SP)* if an agent cannot benefit from misreporting its edges. Formally, strategyproofness requires that for every $i \in N$ and every pair of graphs $G, G' \in \mathcal{G}$ that differ only in the outgoing edges of agent i , it holds that $u_i(G) = u_i(G')$.⁵ This means that the probability of agent $i \in N$ being selected has to be independent of the outgoing edges reported by i . A discussion of this definition in the context of randomized mechanisms can be found in Section 5.

A k -selection mechanism is further called *group strategyproof (GSP)* if there is no coalition of agents that can all gain from jointly misreporting their outgoing edges. Formally, group strategyproofness requires that for every $S \subseteq N$ and every pair of graphs $G, G' \in \mathcal{G}$ that differ only in the outgoing edges of the agents in S , there exists $i \in S$ such that $u_i(G) \leq u_i(G')$. An alternative, stronger definition requires that some agent *strictly* lose as a result of the deviation. Crucially, our result with respect to group strategyproofness is an impossibility, hence using the weaker definition only strengthens the result.

Given a graph G , let $\deg(i) = \deg(i, G)$ be the indegree of agent i in G , i.e., the number of its incoming edges. We seek mechanisms that are SP or GSP, and in addition approximate the optimization target $\sum_{i \in S} \deg(i)$, i.e., we wish to maximize the sum of indegrees of the selected agents. Formally, we say that a k -selection mechanism f provides an *approximation ratio* of α if for every graph G ,

$$\frac{\max_{S \in \mathcal{S}_k} \sum_{i \in S} \deg(i)}{\mathbb{E}_{S \sim f(G)}[\sum_{i \in S} \deg(i)]} \leq \alpha.$$

⁵By symmetry, this is equivalent to writing the last equality as an inequality.

3. DETERMINISTIC MECHANISMS

In this section we study deterministic k -selection mechanisms. Before stating our result, we discuss some special cases.

Clearly, only one mechanism exists if $k = n$, i.e., when all the agents must be selected, and this mechanism is optimal. More interestingly, no deterministic SP mechanism can achieve a finite approximation ratio when $k = 1$. Indeed, let $n \geq 2$, let f be an SP deterministic mechanism, and consider a graph $G = (N, E)$ with $E = \{(1, 2), (2, 1)\}$, i.e., the only two edges are from agent 1 to agent 2 and vice versa. Without loss of generality we may assume that $f(G) = \{1\}$. Now assume that agent 2 removes its outgoing edge, so we obtain the graph $G' = (N, E')$ with $E' = \{(1, 2)\}$. By strategyproofness, $f(G') = \{1\}$, but now agent 2 is the only agent with positive degree, hence the approximation ratio of f is infinite.

Note that in order to achieve a finite approximation ratio, a mechanism must satisfy the following property, which is also sufficient: if there is an edge in the graph, the mechanism must select a subset of agents with at least one incoming edge. The argument above shows that this property cannot be satisfied by any SP mechanism when $k = 1$, but intuitively it should be easy to satisfy when k is large. For the extreme case where $k = n - 1$, for example, the question is whether there exists an SP mechanism with the following very basic property: if there is only one agent with incoming edges, that agent should not be the only one that is *not* selected.

We obtain a surprising negative answer to this question, which turns out to hold even when we restrict our attention to graphs where each agent has at most one outgoing edge. Amusingly, a connection can be made to the popular TV game show “Survivor,” where at the end of each episode the remaining candidates cast votes to choose one of them to be eliminated. Consider a slight variation where each candidate can vote for one other candidate to be eliminated, but is also allowed not to cast a vote. Since each candidate’s first priority is not to be eliminated, strategyproofness in our 0-1 utility model is in fact a necessary condition for strategyproofness in suitable, more refined utility models. The following theorem then implies that a mechanism for choosing the candidate to be eliminated cannot be SP if it is not allowed to eliminate a unique candidate who received any votes. In other words, lies are inherent in the game!

More generally, we show that for *any* value of k , strategyproofness and a finite approximation ratio are mutually exclusive.

THEOREM 3.1. *Let $N = \{1, \dots, n\}$, $n \geq 2$, and $k \in \{1, \dots, n-1\}$. Then there is no deterministic SP k -selection mechanism that provides a finite approximation ratio.*

Proof. Assume for contradiction that $f : \mathcal{G} \rightarrow \mathcal{S}_k$ is a deterministic SP k -selection mechanism with a finite approximation ratio. Furthermore, let $G^* = (N, \emptyset)$ be the empty graph. Since $k < n$, there exists $i \in N$ such that $i \notin f(G^*)$; without loss of generality, $n \notin f(G^*)$.

We now restrict our attention to stars whose center is agent n , i.e., to graphs where the only edges are those of the form (i, n) for an agent $i \in N \setminus \{n\}$. We can represent such a graph by a binary vector $\mathbf{x} = (x_1, \dots, x_{n-1})$, where $x_i = 1$ if and only if the edge (i, n) is in the graph; see Figure 1 for an illustration. In other words, we restrict the domain of f




Figure 1: The star corresponding to the vector $(1, 0, 1, 1, 0, 0)$, for $n = 7$.

to $\{0, 1\}^{n-1}$. Clearly, f must still be SP and provide a finite approximation ratio on the restricted domain.

We claim that $n \in f(\mathbf{x})$ for all $\mathbf{x} \in \{0, 1\}^{n-1} \setminus \{\mathbf{0}\}$. Indeed, in every such graph agent n is the only agent with incoming edges. Hence, any subset that does not include agent n has zero incoming edges, and therefore does not provide a finite approximation ratio.

To summarize, f satisfies the following three constraints:

1. $n \notin f(\mathbf{0})$.
2. For all $\mathbf{x} \in \{0, 1\}^{n-1} \setminus \{\mathbf{0}\}$, $n \in f(\mathbf{x})$.
3. Strategyproofness: for all $i \in N \setminus \{n\}$ and $\mathbf{x} \in \{0, 1\}^{n-1}$, $i \in f(\mathbf{x})$ if and only if $i \in f(\mathbf{x} + e_i)$, where e_i is the i th unit vector and addition is modulo 2.

Next, we claim that $|\{\mathbf{x} \in \{0, 1\}^{n-1} : i \in f(\mathbf{x})\}|$ is even for all $i \in N \setminus \{n\}$. This follows directly from the third constraint, strategyproofness: we can simply partition the set $\{\mathbf{x} \in \{0, 1\}^{n-1} : i \in f(\mathbf{x})\}$ into disjoint pairs of the form $\{\mathbf{x}, \mathbf{x} + e_i\}$.

Finally, we consider the expression $\sum_{\mathbf{x} \in \{0, 1\}^{n-1}} |f(\mathbf{x})|$. On one hand, we have that

$$\begin{aligned} \sum_{\mathbf{x} \in \{0, 1\}^{n-1}} |f(\mathbf{x})| &= \\ \sum_{i \in N} |\{\mathbf{x} \in \{0, 1\}^{n-1} : i \in f(\mathbf{x})\}| &= \quad (1) \\ (2^{n-1} - 1) + \sum_{i \in N \setminus \{n\}} |\{\mathbf{x} \in \{0, 1\}^{n-1} : i \in f(\mathbf{x})\}|, \end{aligned}$$

where the second equality is obtained by separating $|\{\mathbf{x} \in \{0, 1\}^{n-1} : n \in f(\mathbf{x})\}|$ from the sum, and observing that it follows from the first two constraints that this expression equals $2^{n-1} - 1$. Since $2^{n-1} - 1$ is odd and $\sum_{i \in N \setminus \{n\}} |\{\mathbf{x} \in \{0, 1\}^{n-1} : i \in f(\mathbf{x})\}|$ is even, (1) implies that the sum $\sum_{\mathbf{x} \in \{0, 1\}^{n-1}} |f(\mathbf{x})|$ is odd.

On the other hand, it trivially holds that

$$\sum_{\mathbf{x} \in \{0, 1\}^{n-1}} |f(\mathbf{x})| = \sum_{\mathbf{x} \in \{0, 1\}^{n-1}} k = 2^{n-1} \cdot k,$$

hence $\sum_{\mathbf{x} \in \{0, 1\}^{n-1}} |f(\mathbf{x})|$ is even. We have reached a contradiction. \square

It is interesting to note that if we change the problem formulation by allowing the selection of *at most* k agents for $k \geq 2$, then it is possible to design a curious deterministic SP mechanism with a finite approximation ratio. The mechanism first orders the agents lexicographically from left to right. It then scans the agents from left to right, until it

finds an outgoing edge directed to the right, and selects the agent the edge is pointing at. If no edge is found, the last agent to be scanned is selected. The mechanism then does the same for the direction from right to left, again selecting an agent. It is not hard to see that this mechanism is SP and achieves a finite approximation ratio, although this ratio can be as bad as $\Omega(nk)$. Crucially, the mechanism sometimes returns just a single agent, specifically when one and the same agent is selected for the two directions. We refer the reader to Section 5 for further discussion.

4. RANDOMIZED MECHANISMS

In Section 3 we have established a total impossibility result with respect to deterministic SP k -selection mechanisms. In this section we ask to what extent this result can be circumvented using randomization.

4.1 SP Randomized Mechanisms


As we move to the randomized setting, it immediately becomes apparent that Theorem 3.1 no longer applies. Indeed, a randomized SP k -selection mechanism with a finite approximation ratio can be obtained by simply selecting k agents at random. Of course, this mechanism still yields a poor approximation ratio. Can we do better?

Consider first a simple deterministic mechanism that partitions the agents into two predetermined subsets S_1 and S_2 . Next, the mechanism discards all edges between pairs of agents in the same subset. Finally, the mechanism chooses the top $k/2$ agents from each subset. In other words, the mechanism selects the $k/2$ agents with highest indegree from each subset, where the indegree is calculated only on the basis of incoming edges from the other subset. This mechanism is clearly SP. Indeed, consider some $i \in S_t$, $t \in \{1, 2\}$; its outgoing edges to agents inside its subset are disregarded, whereas its outgoing edges to agents in S_{3-t} can only influence which agents are selected from S_{3-t} . However, even without Theorem 3.1, it is easy to see that the mechanism does not yield a finite approximation ratio, since it might be the case that the only edges in the graph are between agents in the same subset.


We will now leverage and refine the partition idea in order to design a randomized SP mechanism that yields a constant approximation ratio. More accurately, we define an infinite family of mechanisms, parameterized by a parameter $m \in \mathbb{N}$. Given m , the mechanism randomly partitions the set of agents into m subsets, and then selects (roughly) the top k/m agents from each subset, based only on the incoming edges from agents in other subsets. Below we give a more formal specification of the mechanism; an example is shown in Figure 2.

The Random m -Partition Mechanism (m -RP)

1. Assign each agent independently and uniformly at random to one of m subsets S_1, \dots, S_m .
2. Let $T \subset \{1, \dots, m\}$ be a random subset of size $k - m \lfloor k/m \rfloor$.
3. If $t \in T$, select the $\lceil k/m \rceil$ agents from S_t with highest indegrees based only on edges from $N \setminus S_t$. If $t \notin T$, select the $\lfloor k/m \rfloor$ agents from S_t with highest indegrees based only on edges from $N \setminus S_t$. Break ties lexicographically in both cases. If one of the subsets S_t is



(a) The given graph



(b) The partitioned graph

Figure 2: Illustration of the Random 2-Partition Mechanism, with $n = 6$ and $k = 2$. Figure 2(a) shows a given graph. The mechanism randomly partitions the agents into two subsets and disregards the edges inside each group, as shown in shown in Figure 2(b). It then selects the best agent in each subset based on the incoming edges from the other subset. In the example, the selected subset is $\{1, 5\}$, with a sum of indegrees of four, whereas the optimal subset is $\{2, 5\}$, with a sum of indegrees of five.

smaller than the number of agents to be selected from this subset, select the entire subset.

4. If only $k' < k$ agents were selected in Step 3, select $k - k'$ additional agents uniformly from the set of agents that were not previously selected.

Note that if $k = 1$ and $m = 2$, we select one agent from one of the two subsets, based on the incoming edges from the other. In this case, step 2 is equivalent to a fair coin toss that determines from which of the two subsets we select an agent.

As in the deterministic case, given a partition of the agents into subsets S_1, \dots, S_m , the choice of agents that are selected from S_t is independent of their outgoing edges. Furthermore, the partition is independent of the input. Therefore, m -RP is SP.⁶ The approximation guarantees provided by m -RP are made explicit in the following theorem.

THEOREM 4.1. *Let $N = \{1, \dots, n\}$, $k \in \{1, \dots, n - 1\}$. For every value of m , m -RP is SP. Furthermore,*

1. 2-RP provides an approximation ratio of four, and
2. for $k \geq 2$, $\left(\lceil k^{1/3} \rceil\right)$ -RP provides an approximation ratio of $1 + \mathcal{O}(1/k^{1/3})$.

⁶The mechanism is even *universally SP*, see Section 5.




Figure 3: An illustration of the proof of Theorem 4.1 for $n = 8$ and $k = 4$. For the given graph, the optimal subset is $K^* = \{1, 2, 3, 4\}$. N is partitioned into $S_1 = \{1, 2, 5, 6\}$ and $S_2 = \{3, 4, 7, 8\}$, which partitions K^* into $K_1^* = \{1, 2\}$ and $K_2^* = \{3, 4\}$. We have that $d_1 = d_2 = 1$.

Proof. For the first part of the theorem, consider an optimal set $K^* \subseteq N$ of k agents (not necessarily unique). Let OPT be the sum of the indegrees of the agents in K^* , i.e.,

$$\text{OPT} = \sum_{i \in K^*} \deg(i).$$

We wish to show that the mechanism selects a k -subset with an expected number of at least $\text{OPT}/4$ incoming edges.

Consider some partition π of the agents into two subsets S_1 and S_2 . In particular, let K^* be partitioned into $K_1^* \subseteq S_1$ and $K_2^* \subseteq S_2$, and assume without loss of generality that $|K_1^*| \geq |K_2^*|$. Denote by d_1 the number of edges from S_2 to K_1^* , i.e.,

$$d_1 = |\{(i, j) \in E : i \in S_2 \wedge j \in K_1^*\}|,$$

and similarly

$$d_2 = |\{(i, j) \in E : i \in S_1 \wedge j \in K_2^*\}|.$$

See Figure 3 for an illustration.

Note that step 2 of the 2-RP Mechanism is equivalent to flipping a fair coin to determine whether we select $\lfloor k/2 \rfloor$ agents from S_1 and $\lfloor k/2 \rfloor$ agents from S_2 (when $T = \{1\}$), or vice versa (when $T = \{2\}$). Now, since $|K_2^*| \leq \lfloor k/2 \rfloor$ (by our assumption that $|K_1^*| \geq |K_2^*|$), it follows that the subset of S_2 selected by the mechanism has at least d_2 incoming edges, regardless of whether $T = \{1\}$ or $T = \{2\}$, and even if $|S_2| < \lfloor k/2 \rfloor$. Moreover, since $|K_1^*| \leq |K^*| = k$, it holds that the subset of S_1 selected by the mechanism has at least $(\lceil k/2 \rceil / k) \cdot d_1$ incoming edges if $T = \{1\}$, and at least $(\lfloor k/2 \rfloor / k) \cdot d_1$ if $T = \{2\}$. Therefore, we have that

$$\begin{aligned} \mathbb{E}[\text{MECH} \mid \pi] &= \\ \mathbb{E}[\text{MECH} \mid \pi \wedge T = \{1\}] &\cdot \frac{1}{2} \\ + \mathbb{E}[\text{MECH} \mid \pi \wedge T = \{2\}] &\cdot \frac{1}{2} \geq \\ \left(\frac{\lfloor k/2 \rfloor}{k} \cdot d_1 + d_2 \right) &\cdot \frac{1}{2} + \left(\frac{\lceil k/2 \rceil}{k} \cdot d_1 + d_2 \right) \cdot \frac{1}{2} = \\ \frac{d_1}{2} + d_2 &\geq \frac{d_1 + d_2}{2}. \end{aligned} \quad (2)$$

For a random partition of the agents into S_1 and S_2 , each edge has probability $1/2$ of being an edge between the two

subsets, and probability $1/2$ of being inside one of the subsets. Hence, by linearity of expectation, the expected number of edges incoming to K^* that are between the two subsets is $\text{OPT}/2$. Formally, for a partition π , let S_1^π and S_2^π be the two subsets of agents, and let

$$d^\pi = |\{(i, j) \in E : (i \in S_1^\pi \wedge j \in S_2^\pi \cap K^*) \vee (i \in S_2^\pi \wedge j \in S_1^\pi \cap K^*)\}|.$$

Then it holds that

$$\sum_{\pi} \Pr[\pi] \cdot d^\pi = \frac{\text{OPT}}{2}. \quad (3)$$

We conclude that

$$\begin{aligned} \mathbb{E}[\text{MECH}] &= \sum_{\pi} \mathbb{E}[\text{MECH} \mid \pi] \cdot \Pr[\pi] \\ &\geq \sum_{\pi} \Pr[\pi] \cdot \frac{d^\pi}{2} = \frac{\text{OPT}}{4}, \end{aligned}$$

where the second transition follows from (2) and the third transition follows from (3).

We now turn to the second part of the theorem. For ease of exposition, and since we are looking for an asymptotic result, we will omit various floors and ceilings from the proof. We employ one additional insight: if k is large enough, the random partition into $k^{1/3}$ subsets will be relatively balanced. A direct approach would be to bound the probability that the number of optimal agents in some subset deviates significantly from $k^{2/3}$, and then proceed in a way similar to the first part. We however take a somewhat different approach that yields a better result.

Consider the agents in the optimal set K^* , and assume without loss of generality that $K^* = \{1, \dots, k\}$. Given $i \in K^*$, we define a random variable Z_i that depends on the random partition of N into $S_1, \dots, S_{k^{1/3}}$ as follows:

$$Z_i = |\{j \in K^* \setminus \{i\} : \exists t \text{ s.t. } i \in S_t \wedge j \in S_t\}|,$$

i.e., Z_i is the number of agents in the optimal set, excluding i itself, that are in the same random subset as agent i . We have

$$\begin{aligned} \mathbb{E}[\text{MECH}] &= \sum_{s_1, \dots, s_k} \mathbb{E}[\text{MECH} \mid Z_1 = s_1, \dots, Z_k = s_k] \\ &\cdot \Pr[Z_1 = s_1, \dots, Z_k = s_k], \end{aligned} \quad (4)$$

where the probability is taken over random partitions.

Recall that the $k^{1/3}$ -RP Mechanism selects the top $k^{2/3}$ agents from each subset. Let

$$\sigma_s = \min\{1, k^{2/3}/(s+1)\}.$$

Furthermore, given $i \in K^*$ and a partition, let

$$d'_i = |\{(j, i) \in E : j \in S_{t_1} \wedge i \in S_{t_2} \wedge t_1 \neq t_2\}|,$$

i.e., d'_i is the number of edges incoming to agent i from other subsets. Using similar arguments to those employed to obtain (2), we get

$$\begin{aligned} \mathbb{E}[\text{MECH} \mid Z_1 = s_1, \dots, Z_k = s_k] &\geq \\ \mathbb{E} \left[\sum_{i \in K^*} d'_i \sigma_{s_i} \mid Z_1 = s_1, \dots, Z_k = s_k \right] &= \\ \sum_{i \in K^*} \mathbb{E}[d'_i \sigma_{s_i} \mid Z_1 = s_1, \dots, Z_k = s_k]. \end{aligned} \quad (5)$$

We wish to obtain an explicit expression for

$$\mathbb{E} [d'_i \sigma_{s_i} \mid Z_1 = s_1, \dots, Z_k = s_k].$$

For $i \in N$ and $S \subseteq N$, let

$$\deg(i, S) = |\{(j, i) \in E : j \in S\}|$$

be the indegree of agent i based on incoming edges from agents in S . We claim that

$$\begin{aligned} \mathbb{E} [d'_i \sigma_{s_i} \mid Z_1 = s_1, \dots, Z_k = s_k] = \\ \left(\frac{k-1-s_i}{k-1} \cdot \deg(i, K^*) + \frac{k^{1/3}-1}{k^{1/3}} \cdot \deg(i, N \setminus K^*) \right) \cdot \sigma_{s_i}. \end{aligned} \quad (6)$$

This identity is obtained by using linearity of expectation twice, as any fixed agent in K^* is not in the same subset as agent i with probability $(k-1-s_i)/(k-1)$, and any fixed agent in $N \setminus K^*$ is not in the same subset as agent i with probability $(k^{1/3}-1)/k^{1/3}$. Notice that the expression on the right hand side of (6) is independent of s_j for all $j \neq i$.

Combining (4), (5), and (6), and reversing the order of summation, we conclude that

$$\begin{aligned} \mathbb{E} [\text{MECH}] &\geq \\ &\sum_{i \in K^*} \sum_{s_1, \dots, s_k} \Pr [Z_1 = s_1, \dots, Z_k = s_k] \cdot \\ &\left(\frac{k-1-s_i}{k-1} \cdot \deg(i, K^*) + \frac{k^{1/3}-1}{k^{1/3}} \cdot \deg(i, N \setminus K^*) \right) \cdot \sigma_{s_i} = \\ &\sum_{i \in K^*} \sum_{s=0}^{k-1} \Pr [Z_i = s] \cdot \\ &\left(\frac{k-1-s}{k-1} \cdot \deg(i, K^*) + \frac{k^{1/3}-1}{k^{1/3}} \cdot \deg(i, N \setminus K^*) \right) \cdot \sigma_s = \\ &\sum_{i \in K^*} \sum_{s=0}^{k-1} \Pr [Z_i = s] \cdot \frac{k-1-s}{k-1} \cdot \deg(i, K^*) \cdot \sigma_s + \\ &\sum_{i \in K^*} \sum_{s=0}^{k-1} \Pr [Z_i = s] \cdot \frac{k^{1/3}-1}{k^{1/3}} \cdot \deg(i, N \setminus K^*) \cdot \sigma_s. \end{aligned}$$

On the other hand, we have that

$$\begin{aligned} \text{OPT} &= \sum_{i \in K^*} (\deg(i, K^*) + \deg(i, N \setminus K^*)) \\ &= \sum_{i \in K^*} \deg(i, K^*) + \sum_{i \in K^*} \deg(i, N \setminus K^*). \end{aligned} \quad (7)$$

In order to complete the proof it therefore suffices to prove that for every $i \in K^*$,

$$\sum_{s=0}^{k-1} \Pr [Z_i = s] \cdot \frac{k^{1/3}-1}{k^{1/3}} \cdot \sigma_s = 1 - \mathcal{O} \left(\frac{1}{k^{1/3}} \right), \quad (8)$$

and

$$\sum_{s=0}^{k-1} \Pr [Z_i = s] \cdot \frac{k-1-s}{k-1} \cdot \sigma_s = 1 - \mathcal{O} \left(\frac{1}{k^{1/3}} \right). \quad (9)$$

Using these equalities we may bound $\mathbb{E} [\text{MECH}]$ from below

by

$$\begin{aligned} &\sum_{i \in K^*} \left(1 - \mathcal{O} \left(\frac{1}{k^{1/3}} \right) \right) \deg(i, K^*) + \\ &\sum_{i \in K^*} \left(1 - \mathcal{O} \left(\frac{1}{k^{1/3}} \right) \right) \deg(i, N \setminus K^*), \end{aligned}$$

and by (7) it follows that

$$\frac{\text{OPT}}{\mathbb{E} [\text{MECH}]} \leq \frac{1}{\left(1 - \mathcal{O} \left(\frac{1}{k^{1/3}} \right) \right)} = 1 + \mathcal{O} \left(\frac{1}{k^{1/3}} \right).$$

Since $\sigma_s = 1$ for all $s \leq k^{2/3} - 1$, in order to establish (8) we must show that

$$\sum_{s=k^{2/3}}^{k-1} \Pr [Z_i = s] \cdot \frac{s+1-k^{2/3}}{s+1} = \mathcal{O} \left(\frac{1}{k^{1/3}} \right).$$

Indeed,

$$\begin{aligned} &\sum_{s=k^{2/3}}^{k-1} \Pr [Z_i = s] \cdot \frac{s+1-k^{2/3}}{s+1} \leq \\ &\sum_{x=1}^{2\sqrt{\log k}} \Pr \left[Z_i \geq k^{2/3} + (x-1)k^{1/3} \right] \cdot \frac{xk^{1/3} + 1}{k^{2/3} + xk^{1/3} + 1} \\ &\quad + \Pr \left[Z_i \geq k^{2/3} + 2\sqrt{\log k} \cdot k^{1/3} \right] \cdot 1. \end{aligned} \quad (10)$$

To bound the probabilities on the right hand side of (10), we employ the following version of the Chernoff bounds (see, e.g., [1], Theorem A.1.11).

LEMMA 4.2. *Let X_1, \dots, X_k be i.i.d. Bernoulli trials where $\Pr [X_i = 1] = p$ for $i = 1, \dots, k$, and define $X = \sum_{i=1}^k X_i$. In addition, let $\lambda > 0$. Then*

$$\Pr [X - kp \geq \lambda] \leq \exp \left(-\frac{\lambda^2}{2kp} + \frac{\lambda^3}{2(kp)^2} \right).$$

Z_i is in fact the sum of $k-1$ i.i.d. Bernoulli trials, but we can safely assume that it is the sum of k trials since we are interested in an upper bound on the probability that the sum is greater than a certain value. Using Lemma 4.2 with $\lambda = xk^{1/3}$ and $p = 1/k^{1/3}$, we get

$$\begin{aligned} &\Pr \left[Z_i \geq k^{2/3} + (x-1)k^{1/3} \right] \leq \\ &\exp \left(-\frac{(x-1)^2 k^{2/3}}{2k^{2/3}} + \frac{(x-1)^3 k}{2k^{4/3}} \right) \leq \exp \left(-\frac{(x-1)^2}{4} \right), \end{aligned} \quad (11)$$

where the second inequality holds for a large enough k . Similarly,

$$\begin{aligned} &\Pr \left[Z_i \geq k^{2/3} + 2\sqrt{\log k} \cdot k^{1/3} \right] \leq \\ &\exp \left(-\frac{4k^{2/3} \log k}{2k^{2/3}} + \frac{8k(\log k)^{3/2}}{2k^{4/3}} \right) \leq \exp(-\log k) \leq \frac{1}{k}. \end{aligned}$$

We conclude that the expression on the right hand side

of (10) is bounded from above by

$$\sum_{x=1}^{2\sqrt{\log k}} \left(\exp\left(-\frac{(x-1)^2}{4}\right) \cdot \frac{xk^{1/3} + 1}{k^{2/3} + xk^{1/3} + 1} \right) + \frac{1}{k} \leq \frac{1}{k^{1/3}} \sum_{x=1}^{2\sqrt{\log k}} \left(\exp\left(-\frac{(x-1)^2}{4}\right) \cdot 2x \right) + \frac{1}{k} = \mathcal{O}\left(\frac{1}{k^{1/3}}\right),$$

which follows from the fact that the series $\sum_{x=1}^{\infty} \exp(-\Theta(x^2)) \cdot \Theta(x)$ converges. This establishes (8).

The proof of (9) is similar to that of (8). It is sufficient to show that

$$\sum_{s=0}^{k^{2/3}-1} \Pr[Z_i = s] \cdot \frac{s}{k-1} + \sum_{s=k^{2/3}}^{k^{2/3}+2\sqrt{\log k} \cdot k^{1/3}-1} \Pr[Z_i = s] \left(1 - \frac{k-1-s}{k-1} \cdot \frac{k^{2/3}}{s+1}\right) + \Pr\left[Z_i \geq k^{2/3} + 2\sqrt{\log k} \cdot k^{1/3}\right] \cdot 1 = \mathcal{O}\left(\frac{1}{k^{1/3}}\right).$$

It holds that

$$\sum_{s=0}^{k^{2/3}-1} \Pr[Z_i = s] \cdot \frac{s}{k-1} \leq \sum_{s=0}^{k^{2/3}-1} \Pr[Z_i = s] \cdot \frac{k^{2/3}-1}{k-1} = \mathcal{O}\left(\frac{1}{k^{1/3}}\right),$$

and as before,

$$\Pr\left[Z_i \geq k^{2/3} + 2\sqrt{\log k} \cdot k^{1/3}\right] \cdot 1 \leq \frac{1}{k}.$$

Finally,

$$\sum_{s=k^{2/3}}^{k^{2/3}+2\sqrt{\log k} \cdot k^{1/3}-1} \Pr[Z_i = s] \left(1 - \frac{k-1-s}{k-1} \cdot \frac{k^{2/3}}{s+1}\right) = \sum_{s=k^{2/3}}^{k^{2/3}+2\sqrt{\log k} \cdot k^{1/3}-1} \Pr[Z_i = s] \cdot \left(1 - \left(1 - \mathcal{O}\left(\frac{1}{k^{1/3}}\right)\right) \cdot \frac{k^{2/3}}{s+1}\right).$$

We can thus bound this sum from above as before using (11). This completes the proof. \square

When using m -RP, we can in fact choose the best value of m for any given value of k . In other words, Theorem 4.1 implies that for every k there exists an SP mechanism with an approximation ratio of $\min\{4, 1 + \mathcal{O}(1/k^{1/3})\}$, i.e., an approximation ratio that is bounded from above by four for any value of k , and approaches one as k grows.

For $k = 1$, 2-RP provides an approximation ratio of four, while that of m -RP with $m > 2$ is strictly worse. It is interesting to note that the analysis in this case is tight. To see this, consider a graph $G = (N, E)$ with only one edge from agent 1 to agent n , i.e., $E = \{(1, n)\}$. Assume without loss of generality that agent n is assigned to S_1 . We distinguish two cases, depending on the value of T :

1. If $T = \{1\}$, then agent n is selected if and only if $1 \in S_2$. Indeed, if $1 \in S_2$, then agent n is the only agent in S_1 with a positive indegree from S_2 , and will

therefore be selected. If instead $1 \in S_1$, then all agents in S_1 have indegree zero from S_2 , and agent 1 will be chosen due to lexicographic tie-breaking. Conditioned on $T = \{1\}$, n is selected with probability $1/2$.

2. If $T = \{2\}$, then agent n is selected if and only if $S_2 = \emptyset$ and if n is chosen from a uniform distribution over $S_1 = N$. Conditioned on $T = \{2\}$, n is selected with probability $(1/2^{n-1}) \cdot (1/n)$.

Since the value of T is determined by a fair coin toss, the probability that n is selected by 2-RP is therefore exactly $1/4 + 1/(2^n \cdot n)$. We conclude that the approximation ratio of 2-RP cannot be smaller than

$$\frac{1}{\frac{1}{4} + \frac{1}{2^n \cdot n}} = 4 - \mathcal{O}\left(\frac{1}{2^n \cdot n}\right).$$

We next provide a very simple, though rather weak, lower bound for the approximation ratio achievable by any randomized SP k -selection mechanism. Let $k \in \{1, \dots, n-1\}$, and let $f : \mathcal{G} \rightarrow \Delta(\mathcal{S}_k)$ be a randomized SP k -selection mechanism. Consider the graph $G = (N, E)$ where

$$E = \{(i, i+1) : i = 1, \dots, k\} \cup \{(k+1, 1)\},$$

i.e., E is a directed cycle on the agents $1, \dots, k+1$. Then there exists an agent $i \in \{1, \dots, k+1\}$, without loss of generality agent 1, that is included in $f(G)$ with probability at most $k/(k+1)$. Now consider the graph $G' = (N, E')$ with $E' = E \setminus \{(1, 2)\}$, which is obtained from G if agent 1 removes its outgoing edge to agent 2. By strategyproofness, agent 1 is included in $f(G')$ with probability at most $k/(k+1)$. Any subset $S \in \mathcal{S}_k$ such that $1 \notin S$ has at most $k-1$ incoming edges in G' . It follows that the expected number of incoming edges in $f(G')$ is at most

$$\frac{k}{k+1} \cdot k + \frac{1}{k+1} \cdot (k-1) = \frac{k^2 + k - 1}{k+1},$$

so the approximation ratio of f cannot be smaller than

$$\frac{k}{\frac{k^2+k-1}{k+1}} = 1 + \frac{1}{k^2 + k - 1}. \quad (12)$$

We have thus proved the following easy result.

THEOREM 4.3. *Let $N = \{1, \dots, n\}$, $n \geq 2$, $k \in \{1, \dots, n-1\}$. Then there is no randomized SP k -selection mechanism with an approximation ratio smaller than $1 + \Omega(1/k^2)$.*

Not surprisingly, this lower bound converges to one, albeit more quickly than the upper bound of Theorem 4.1. Again, the special case where $k = 1$ is particularly interesting: here, (12) yields an explicit lower bound of two, while Theorem 4.1 provides an upper bound of four. We conjecture that the correct value is two.

CONJECTURE 4.4. *There exists a randomized SP 1-selection mechanism with an approximation ratio of two.*

One deceptively promising avenue for proving the conjecture is to design an iterative version of the Random Partition Mechanism. In particular, we could start with an empty subset $S \subset N$, and in each step add to S an agent from $N \setminus S$ that has minimum indegree based on the incoming edges from S , breaking ties randomly (so in the first step we would just add a random agent). The last agent that remains outside S would then be selected. This mechanism is SP and does remarkably well on some difficult instances, but fails spectacularly on a contrived counterexample.

4.2 GSP Randomized Mechanisms

In the beginning of Section 4.1 we identified a trivial randomized SP k -selection mechanism, namely the one that selects a subset of k agents at random. Of course this mechanism is even GSP, since the outcome is completely independent of the reported graph.

We claim that selecting a random k -subset provides an approximation ratio of n/k . Indeed, consider an optimal subset $K^* \subseteq N$ with $|K^*| = k$. Each agent $i \in K^*$ is included in the selected subset with probability k/n , and hence in expectation contributes a (k/n) -fraction of its indegree to the expected total indegree of the selected subset. By linearity of expectation, the expected total indegree of the selected subset is at least a (k/n) -fraction of the total indegree of K^* .

Theorem 4.1 implies that we can do much better if we just ask for strategyproofness. If one asks for *group* strategyproofness, on the other hand, selecting a random subset turns out to be optimal up to a tiny gap.

THEOREM 4.5. *Let $N = \{1, \dots, n\}$, $n \geq 2$, and let $k \in \{1, \dots, n-1\}$. No randomized GSP k -selection mechanism can provide an approximation ratio smaller than $(n-1)/k$.*

Proof. Let $f : \mathcal{G} \rightarrow \mathcal{S}_k$ be a randomized GSP mechanism. Let $G = (N, \emptyset)$ be the empty graph, and observe that there must exist two agents $i, j \in N$ such that $f(G)$ selects each of them with probability at most $k/(n-1)$.

Now consider the graph $G' = (N, E')$ with $E' = \{(i, j), (j, i)\}$. By group strategyproofness, it must hold for either i or j that the probability that this agent is selected by $f(G')$ is no greater than the probability that this agent is selected by $f(G)$; we may thus assume without loss of generality that $f(G')$ selects i with probability at most $k/(n-1)$.

Now consider the graph G'' with $E'' = \{(j, i)\}$. By strategyproofness, i is selected with equal probability in $f(G')$ and $f(G'')$, i.e., with probability at most $k/(n-1)$. Since i is the only agent with an incoming edge in G'' , the approximation ratio is at least $(n-1)/k$. \square

Note that this result holds even if one is merely interested in coalitions of size at most two.

5. DISCUSSION

In this section we discuss the significance of our results and state some open problems.

Payments If payments are allowed and the preferences of the agents are quasi-linear, truthful implementation of the optimal solution is straightforward: simply give one unit of payment to each agent that is not selected. This can be refined by only paying “pivotal” agents that are not selected, i.e., agents that would have been selected had they misreported their preferences. However, even under the latter scheme we may have to pay all non-selected agents, e.g., when the graph is a clique. A simple argument shows that no truthful payment scheme can do better.

The utility model We have studied an “extreme” utility model, where an agent is only interested in the question of its own selection. This restriction of the preferences of the agents allows us to circumvent impossibility results that hold with respect to more general preferences, e.g., the Gibbard-Satterthwaite Theorem [13, 21] and its generalization to randomized rules [14].

A more practical assumption would be that an agent receives a utility of one if it is selected, plus a utility of $\beta \geq 0$ for each of its (outgoing) neighbors that is selected. In this case, selecting a set S of agents yields social welfare (i.e., sum of utilities) k plus β times the total indegree of S . Hence, if $\beta > 0$, a set S maximizes social welfare if and only if it maximizes the total indegree. In particular, if $\beta > 0$ and payments are available, we can use the VCG mechanism (see, e.g., [18]) to maximize the total indegree in a truthful way.

It is easy to see that the lower bound of Theorem 3.1 for the 0–1 model also holds for the β -1 model if β is small. The latter is likely to be the case in many practical settings, such as those described in Section 1. Upper bounds identical to those of Theorem 4.1 hold for any value of β . In particular, m -RP remains strategyproof in the β -1 model, as the probability that an agent is selected increases in the number of votes it receives. Moreover, if β is small, a variation of the random partition mechanism achieves an approximation ratio close to one with respect to social welfare, even when $k = 1$. If $\beta \geq 1$ then simply selecting the optimal solution (and breaking ties lexicographically) is SP.

Robustness of the impossibility result Theorem 3.1 provides a strong impossibility result for deterministic mechanisms. We have seen that this result is rather sensitive to the model, and no longer holds if one is allowed to select *at most* k agents rather than exactly k , or if each agent is forced to report at least one outgoing edge. That said, we note that these particular aspects of the model are crucial: in our motivating examples, and in approval voting in general, an agent may choose not to report any outgoing edges; in essentially all conceivable applications the set of agents to be selected is of fixed size.

Weights and an application to conference reviews

A seemingly natural generalization of our model can be obtained by allowing weighted edges. Interestingly, our main positive result, Theorem 4.1, also holds in this more general setting (subject to minor modifications of its formulation and proof). Closer scrutiny reveals, however, that our target function is often meaningless in the weighted setting. Indeed, the absence of an edge between i and j would in this context imply that i has no information about j , whereas an edge with small weight would imply that i dislikes or distrusts j . Therefore, maximizing the sum of weights on incoming edges may not be desirable.

That said, in very specific situations maximizing the sum of weights on incoming edges makes perfect sense; one prominent example are conference reviews. In this context, a subset of papers must be selected based on scores assigned by reviewers, who often also submit papers of their own. What is special about this setting is that each paper is usually reviewed by a fixed number of reviewers (say, three), so each agent has the same number of incoming weighted edges, and maximizing the sum of scores is the same as maximizing the average score. This means that m -RP can be used to build a conference program in a truthful and approximately optimal way!

Universal strategyproofness vs. strategyproofness in expectation

In the context of randomized mechanisms, two flavors of strategyproofness are usually considered. A mechanism is *universally SP* if for every fixed outcome of the random choices made by the mechanism an agent cannot gain by lying, i.e., the mechanism is a distribution over

SP mechanisms. A mechanism is *SP in expectation* if an agent cannot increase its expected utility by lying. In this paper we have used the latter definition, which clearly is the weaker of the two. On the one hand, this strengthens the randomized SP lower bound of Theorem 4.3. On the other hand, the randomized mechanisms of Section 4 are in fact universally SP. Indeed, for every fixed partition, selecting agents from one subset based on incoming edges from other subsets is SP. This makes Theorem 4.1 even stronger than originally stated.

Open problems Our most enigmatic open problem is the gap for randomized SP 1-selection mechanisms: Theorem 4.1 yields an upper bound of four, while Theorem 4.3 yields a lower bound of two. We conjecture that there exists a randomized SP 1-selection mechanism that provides a 2-approximation.

A potentially interesting variation of our problem can be obtained by changing the target function. One attractive option is to maximize the minimum indegree in the selected subset. Clearly, our total impossibility regarding deterministic SP mechanisms (Theorem 3.1) carries over to this new target function. However, it is unclear what can be achieved using randomized SP mechanisms.

6. ACKNOWLEDGMENTS

We thank Moshe Babaioff, Liad Blumrosen, Michal Feldman, Ron Holzman, Gil Kalai, Hervé Moulin, David Parkes, Yoav Shoham, and Aviv Zohar for valuable discussions. Alon is supported in part by a USA Israeli BSF grant, by a grant from the Israel Science Foundation, by an ERC advanced grant, and by the Hermann Minkowski Minerva Center for Geometry at Tel Aviv University. Fischer is supported by DFG grants BR 2312/3-2 and FI 1664/1-1 and by a Minerva Short-Term Research Grant. Part of the work was done while Fischer was visiting the Hebrew University of Jerusalem, and Procaccia was at Microsoft Israel R&D Center.

REFERENCES

- [1] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley, 2nd edition, 2000.
- [2] A. Altman and M. Tennenholtz. Strategyproof deterministic lotteries under broadcast communication. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 1549–1552, 2008.
- [3] A. Altman and M. Tennenholtz. Axiomatic foundations for ranking systems. *Journal of Artificial Intelligence Research*, 31:473–495, 2008.
- [4] A. Altman, A. D. Procaccia, and M. Tennenholtz. Non-manipulable selections from a tournament. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 27–32, 2009.
- [5] S. J. Brams and P. C. Fishburn. *Approval Voting*. Springer, 2nd edition, 2007.
- [6] I. Caragiannis, D. Kalaitzis, and E. Markakis. Approximation algorithms and mechanism design for minimax approval voting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 737–742, 2010.
- [7] A. Cheng and E. Friedman. Sybilproof reputation mechanisms. In *Proceedings of the 3rd Workshop on the Economics of Peer-to-Peer Systems*, pages 128–132, 2005.
- [8] O. Dekel, F. Fischer, and A. D. Procaccia. Incentive compatible regression learning. *Journal of Computer and System Sciences*, 76(8):759–777, 2010.
- [9] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 644–652, 2006.
- [10] U. Feige, A. Flaxman, J. D. Hartline, and R. Kleinberg. On the competitive ratio of the random sampling auction. In *Proceedings of the 1st International Workshop on Internet and Network Economics*, pages 878–886, 2005.
- [11] A. Fiat, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Competitive generalized auctions. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 72–81, 2002.
- [12] E. Friedman, P. Resnick, and R. Sami. Manipulation-resistant reputation systems. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 27. Cambridge University Press, 2007.
- [13] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.
- [14] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45:665–681, 1977.
- [15] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55:242–269, 2006.
- [16] R. Holzman and H. Moulin. Impartial award of a prize. Manuscript, 2010.
- [17] R. LeGrand, E. Markakis, and A. Mehta. Some results on approximating the minimax solution in approval voting. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 1193–1195, 2007.
- [18] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9. Cambridge University Press, 2007.
- [19] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1–2):166–196, 2001.
- [20] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 177–186, 2009.
- [21] M. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [22] J. Schummer and R. V. Vohra. Mechanism design without money. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 10. Cambridge University Press, 2007.