# Sum of Us: Truthful Self-Selection

Noga Alon[*]    Felix Fischer[†]    Ariel D. Procaccia[‡]    Moshe Tennenholtz [§]

## Abstract

We consider directed graphs over a set of $n$ agents, where an edge $(i, j)$ is taken to mean that agent $i$ supports or trusts agent $j$. Given such a graph and an integer $k \leq n$, we wish to select a subset of $k$ agents that maximizes the sum of indegrees, i.e., a subset of $k$ most popular or most trusted agents. At the same time we assume that each individual agent is only interested in being selected, and may misreport its outgoing edges to this end. This problem formulation captures realistic scenarios where agents choose among themselves, which can be found in the context of Internet search, social networks like Twitter, or reputation systems like Epinions.

Our goal is to design mechanisms without payments that map each graph to a $k$-subset of agents to be selected and satisfy the following two constraints: *strategyproofness*, i.e., agents cannot benefit from misreporting their outgoing edges, and *approximate optimality*, i.e., the sum of indegrees of the selected subset of agents is always close to optimal. Our first main result is a surprising impossibility: for $k \in \{1, \ldots, n-1\}$, no deterministic strategyproof mechanism can provide a finite approximation ratio. Our second main result is a randomized strategyproof mechanism with an approximation ratio that is bounded from above by four for any value of $k$, and approaches one as $k$ grows.

# 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 that it approves; this setting is known as *Approval voting* [6].

We consider the special case of Approval voting when the set of agents and the set of alternatives coincide. Specifically, in our model there is an underlying directed graph, with the agents as vertices. An edge from agent $i$ to agent $j$ implies that agent $i$ approves, votes for, trusts, or supports agent $j$. Our goal is to select a subset of $k$ "best" agents, based on the given graph; we elaborate on what we mean by "best" momentarily.

Our assumption that agents and alternatives coincide enables us to restrict the agents' preferences. Indeed, we assume that each agent is only interested in whether it is among those selected, that is, it receives utility one if selected and zero otherwise. This assumption reflects, in the limit, a situation where each agent gives very small weight to the overall composition of the selected subset, and very high weight to the question of its own selection.[1]

As a first motivating example, consider an Internet search setting. The web sites are the agents, while hyperlinks are represented by edges. Given this graph, a search engine must return a set of the, say, ten top web sites. Put another way, the top web sites are selected based on the votes cast by other web sites in the form of hyperlinks. Each specific web site or, more accurately, its webmaster is naturally concerned only with appearing at the top of the search results, and to this end may add or remove hyperlinks at will.

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 outgoing edges in the underlying graph $G$ are private information of the respective agent. Fixing a mechanism $f$, the agents play the following game. Each of them reports to the mechanism a set of outgoing edges, which might differ from the true set. The reported edges induce a graph $G'$, and the mechanism selects the subset $f(G')$. We say that a mechanism is *strategyproof (SP)* if an agent cannot benefit from misreporting its outgoing edges, that is, cannot increase its chances of being selected, 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.

We now explain 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. The goal of the mechanism designer is to optimize this target function. Note that this goal is in a sense orthogonal to the agent's interests, which may make the design of good SP mechanisms difficult.

A second motivating example can be found in the context of social networks. While some social networks, like Facebook (`http://facebook.com`), correspond to undirected graphs, there are also many examples with unilateral connections. Each user of the reputation system Epinions (`http://epinions.com`) has a "Web of Trust", that is, the user unilaterally chooses which other users to trust. Another prominent example is the social network Twitter (`http://twitter.com`),

---

[1]See Section 5 for further discussion of this utility model.

which of late has become wildly popular; a Twitter user may choose which other users to "follow."

In "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 pure fantasy, but, indeed, celebrity users of Twitter have recently 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 being selected by not following any other users, that is, reporting an empty set of outgoing edges.

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 give 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 clearly 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 (Procaccia and Tennenholtz [22]), building on earlier work (for example by Dekel et al. [10]). 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* [21] considers mechanisms that are allowed to transfer payments to and from the agents, money is usually unavailable in Internet domains like the ones discussed above (social networks, search engines) due to security and accountability issues (see, e.g., the book chapter by Schummer and Vohra [24]). Our notion of a mechanism, sometimes referred to as a *social choice rule* in the social choice literature, therefore precludes payments by definition. Note that Procaccia and Tennenholtz [22], and also Alon et al. [1], deal with a completely different domain, namely facility location.

For $k = 1$, that is, if one agent must be selected, the game we deal with is a special case of so-called *selection games* [4], where the possible strategies are the outgoing edges. More generally, this setting is related to work in distributed computing on *leader election* (see, e.g., [2, 9, 12, 5]). This line of work does not deal with self-interested agents. Instead, there is a certain number of malicious agents trying to manipulate the selection process, and the goal is to guarantee the selection of a non-malicious agent, at least with a certain probability.

Finally, this paper is related to work on manipulation of reputation systems, which are often modeled as weighted directed graphs; a reputation function maps a given graph to reputation values for the agents (see, e.g., [7, 15]). Although 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 reasonable target function. Indeed, in this context the absence of a specific incoming edge (which indicates lack of knowledge) is preferable to an edge with low weight (which indicates distrust).

**Our 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. A summary of our results is given in Table 1.

Our contribution begins in Section 3 with a study of deterministic $k$-selection mechanisms. It is quite easy to see that no deterministic SP 1-selection mechanism can yield a finite approximation

|  |  | Deterministic | Randomized |
|---|---|:---:|:---:|
| **SP** | **Upper bound** | n/a | $\min\{4, 1 + \mathcal{O}(1/k^{1/3})\}$ |
|  | **Lower bound** | $\infty$ | $1 + \Omega(1/k^2)$ |
| **GSP** | **Upper bound** | n/a | $\frac{n}{k}$ |
|  | **Lower bound** | $\infty$ | $\frac{n-1}{k}$ |

Table 1: Summary of our results for $k$-selection mechanisms, where $n$ is the number of agents. SP stands for strategyproof, GSP for group strategyproof.

ratio. Intuitively, this should not be true for large values of $k$. Indeed, in order to have 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 the 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 $N = \{1, \ldots, n\}$, $n \geq 2$, and $k \in \{1, \ldots, n-1\}$. Then there is no deterministic SP $k$-selection mechanism that gives a finite approximation ratio.*

The proof of the theorem is compact 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 over 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, *Random m-Partition (m-RP)*, 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 [14, 18, 13] and combinatorial auctions [11], although our problem is fundamentally different. We have the following theorem.

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

1. *2-RP has an approximation ratio of four, and*

2. *$\left(\lceil k^{1/3} \rceil\right)$-RP has 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. Thus, there exists a mechanism that always yields an approximation ratio of at most four, and furthermore provides a ratio that approaches one as $k$ grows. In addition, we prove a lower bound of $1 + \Omega(1/k^2)$ on the approximation ratio that can be achieved 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 gives an approximation ratio of $n/k$.

3

## 2 The Model

Let $N = \{1, \ldots, n\}$ be a set of *agents*. For each $k = 1, \ldots, n$, let $\mathcal{S}_k = \mathcal{S}_k(n)$ be the collection of $k$-subsets of $N$, i.e., $\mathcal{S}_k = \{S \subseteq N : |S| = k\}$. We consider directed graphs $G = (N, E)$, that is, graphs with $N$ as the set of vertices, and write $\mathcal{G} = \mathcal{G}(N)$ for the set of such graphs.

A *deterministic $k$-selection mechanism* is a function $f : \mathcal{G} \to \mathcal{S}_k$ that selects a subset of agents for each graph. 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 $f : \mathcal{G} \to \Delta(\mathcal{S}_k)$, 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 can be seen as a special case of a randomized ones, always selecting a set of agents 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 $u_i(G) = u_i(G')$.[2] 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 *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 $\sum_{i \in S} \deg(i)$, that is, we wish to maximize the sum of indegrees of the selected agents. Formally, we say that a $k$-selection mechanism $f$ has an approximation ratio of $\alpha$ if for every graph $G$,

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

## 3 Deterministic Mechanisms

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

Clearly, only one mechanism exists for $k = n$, that is, when all the agents must be selected, and this mechanism is optimal. More interestingly, it is easy to see that one cannot obtain a finite approximation ratio via a deterministic SP mechanism 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.,

---

[2]By symmetry, this is equivalent to writing the last equality as an inequality.

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; formally, we now consider 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 have a finite approximation ratio, our 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 very large.

Consider, for example, the case where $k = n - 1$, that is, the mechanism must select all the agents save one. Can we design an SP mechanism with the extremely basic property that if there is only one agent with incoming edges, that agent would not be the only one *not* to be selected?

In the following theorem, we give a surprising negative answer to this question, even when we restrict our attention to graphs where each agent has at most one outgoing edge. Amusingly, a connection to the popular TV game show "Survivor" can be made. Consider a slight variation where each tribe member can vote for one other trusted member, but is also allowed not to cast a vote. One member must be eliminated at the tribal council, based on the votes. Since each member's first priority is not to be eliminated (i.e., to be selected), strategyproofness in our 0–1 utility model is in fact a necessary condition for strategyproofness in suitable, more refined utility models. The theorem then implies that a mechanism for choosing the eliminated member cannot be SP (even under 0–1 utilities) if it has the property that a member who is the only one that received votes cannot be eliminated. Put another way, lies are inherent in the game!

More generally, we show that for *any* value of $k$, strategyproofness and finite approximation ratio are mutually exclusive. The proof is concise but nontrivial.

**Theorem 3.1.** *Let $N = \{1, \ldots, n\}$, $n \geq 2$, and $k \in \{1, \ldots, n-1\}$. There is no deterministic SP $k$-selection mechanism that gives a finite approximation ratio.*

*Proof.* Assume for contradiction that $f : \mathcal{G} \to \mathcal{S}_k$ is a deterministic SP $k$-selection mechanism that gives 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 will restrict our attention to stars whose center is agent $n$, that is, graphs where the only edges are of the form $(i, n)$ for an agent $i \in N \setminus \{n\}$. We can represent such a graph by a binary vector $\mathbf{x} = (x_1, \ldots, x_{n-1})$, where $x_i = 1$ if and only if the edge $(i, 1)$ is in the graph; see Figure 1 for an illustration. In other words, we restrict the domain of $f$ to $\{0, 1\}^{n-1}$.

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 give a finite approximation ratio (as a subset that does include agent $n$ has at least one incoming edge).

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.

5

(a) The vector $(1,0,1,1,0,0)$      (b) The vector $(1,1,0,0,0,1)$

Figure 1: Correspondence between stars and binary $(n-1)$-vectors, for $n = 7$

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

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

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 $\sum_{\mathbf{x} \in \{0,1\}^{n-1}} |f(\mathbf{x})|$ 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. $\qquad \square$

It is interesting to note that if we slightly 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 that selects at most two agents. The reader is referred to Appendix B for more details.

## 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

(a) The given graph

(b) The partitioned graph

Figure 2: Example for the Random 2-Partition Mechanism, with $n = 6$ and $k = 2$. Figure 2(a) illustrates the given graph. The mechanism randomly partitions the agents into two subsets, shown in Figure 2(b), and disregards the edges inside each group. The mechanism then selects the best agent in each group based on the incoming edges from the other group; 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.

can be obtained by simply selecting $k$ agents at random. However, this mechanism still yields a poor approximation ratio. Can we do better?

The main result of this section is 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 can be found in Figure 2.

**The Random $m$-Partition Mechanism ($m$-RP)**

1. Assign each agent independently and uniformly at random to one of $m$ subsets $S_1, \ldots, S_m$.

2. Let $T \subset \{1, \ldots, m\}$ be a random subset of size $k - m \cdot \lfloor k/m \rfloor$.

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 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 $m = 2$, then 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 toss of a fair coin that determines from which of the two subsets we select an agent.

Given a partition of the agents into subsets $S_1, \ldots, S_m$, the choice of agents that are selected from $S_t$ is completely independent of the outgoing edges of the agents in $S_t$. Furthermore, the

7

partition is independent of the input. Therefore, $m$-RP is SP.[3] The following theorem explicitly states the approximation guarantees provided by $m$-RP; the technical and rather delicate proof of the theorem is relegated to Appendix A.

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

1. *2-RP has an approximation ratio of four, and*

2. *$\left(\lceil k^{1/3} \rceil\right)$-RP has an approximation ratio of $1 + \mathcal{O}(1/k^{1/3})$.*

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

It follows from the theorem that, for $k = 1$, 2-RP has an approximation ratio of four; for this case $m$-RP with $m > 2$ has a strictly worse ratio. It is interesting to note that the analysis is tight. Indeed, consider a graph $G = (N, E)$ with only one edge from agent 1 to agent $n$, that is, $E = \{(1, n)\}$. Assume without loss of generality that agent $n$ is assigned to $S_1$. In order for agent $n$ to be selected, two events must occur:

1. $T = \{1\}$, that is, the winner must be selected from $S_1$. This happens with probability $1/2$.

2. Either $1 \in S_2$, or $|S_1| = 1$. The probability that $1 \in S_2$ is $1/2$. The probability that $|S_1| = 1$, given that $n \in S_1$, is $1/2^{n-1}$. By the union bound, the probability of this event is at most $1/2 + 1/2^{n-1}$.

It is clear that $n$ cannot be selected unless the first event occurs. If the second event does not occur, it follows that $n$ has an indegree of zero based on the incoming edges from $S_2$, and there are other alternatives in $S_1$ (which also have an indegree of zero). Since tie-breaking is lexicographic, agent $n$ would not be selected. As the two events are independent, the probability of both occurring is therefore at most $1/4 + 1/2^n$. We conclude that the approximation ratio of the mechanism cannot be smaller than

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

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

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

i.e., $E$ is a directed cycle on the agents $1, \ldots, k+1$. Then there exists an agent $i \in \{1, \ldots, k+1\}$, without loss of generality agent 1, that is included in $f(G)$ with probability at most $k/(k+1)$. Now, consider the graph $G'$ where $E' = E \setminus \{(1, 2)\}$, that is, agent 1 removes its outgoing edge to agent 2. By strategyproofness, agent 1 is included in $f(G')$ with probability at most $k/(k+1)$.

---

[3]The mechanism is even *universally SP*, see Section 5.

8

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

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

Hence the approximation ratio of $f$ cannot be smaller than

$$\frac{k}{\frac{k^2+k-1}{k+1}} = 1 + \frac{1}{k^2 + k - 1}. \tag{2}$$

We have therefore proved the following easy result.

**Theorem 4.2.** *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, the lower bound given by Theorem 4.2 converges to one, albeit more quickly than the upper bound of Theorem 4.1. As usual, an especially interesting special case is when $k = 1$. Equation (2) gives an explicit lower bound of two for this case. On the other hand, Theorem 4.1 gives an upper bound of four. We conjecture that the correct value is two.

**Conjecture 4.3.** *There exists a randomized SP 1-selection mechanism with an approximation ratio of two.*

One deceptively promising avenue for proving the conjecture is designing an iterative version of the Random Partition Mechanism. Specifically, we start with an empty subset $S \subset N$, and at 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 to $S$ a random agent). The last agent that remains outside $S$ is selected. This SP mechanism does remarkably well on some difficult instances, but fails spectacularly on a contrived counterexample. We give a formal specification of this *Sliding Partition Mechanism*, and construct the illuminating counterexample, in Appendix C.

## 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 gives 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, just selecting a random subset turns out to be optimal up to a tiny gap.

**Theorem 4.4.** *Let* $N = \{1, \dots, n\}$, $n \geq 2$, *and let* $k \in \{1, \dots, n-1\}$. *No randomized GSP $k$-selection mechanism can yield an approximation ratio smaller than $(n-1)/k$.*

*Proof.* Let $f : \mathcal{G} \to \mathcal{S}_k$ be a randomized GSP mechanism. Given the empty graph, there are two agents $i, j \in N$ such that each is selected with probability at most $k/(n-1)$.

Consider the graph $G'$ where $E' = \{(i, j), (j, i)\}$, that is, there are only two edges in $G'$, from $i$ to $j$ and from $j$ to $i$. By group strategyproofness, it must hold for either $i$ or $j$ that $f(G')$ selects this agent with probability not greater than under the empty graph; we may 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 under $f(G')$ and $f(G'')$, that is, 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 $(n-1)/k$. $\qquad\square$

Note that Theorem 4.4 holds even if one is merely interested in coalitions of size at most two.

# 5   Discussion

In this section we discuss three of the issues mentioned above, and list some open problems.

**Payments.** If payments are allowed and the preferences of the agents are quasi-linear, then 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, that is, agents that would have been selected had they lied. However, even under the latter scheme we may have to pay all the non-selected agents (e.g., when the graph is a clique). Moreover, a simple argument shows that there is no truthful payment scheme that does better.

**The utility model.** We have studied an "extreme" utility model, where an agent is only interested in the question of its own selection. The 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 [16, 23] and its generalization to randomized rules [17].

It is possible to consider a more sensitive utility function, where 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 model the social welfare (sum of utilities) of a set $S$ of selected agents is $k$ plus $\beta$ times the total indegree of $S$. Hence, if $\beta > 0$, a set $S$ maximizes the 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 [25, 8, 19] (see [20] for an overview) to maximize the total indegree in a truthful way.

It is easy to verify that any upper bound in the 0–1 model (with total indegree as the target function) also holds in the $\beta$–1 model (with social welfare as the target function), hence Theorem 4.1 is true for the latter. Furthermore, if not zero, $\beta$ may still be very small in many settings, like those described in Section 1. In such cases a variation on the random partition mechanism achieves an approximation ratio close to one for the social welfare, even when $k = 1$. Finally, note that if $\beta \geq 1$ then simply selecting the optimal solution (and breaking ties lexicographically) is SP.

**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, that is, 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 4.2. On the other hand, notice that 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. Hence, Theorem 4.1 is even stronger than originally stated.

**Open problems.** Our most enigmatic open problem is the gap for randomized SP 1-selection mechanisms: Theorem 4.1 gives an upper bound of four, while Theorem 4.2 gives a lower bound of two. We conjecture that there exists a randomized SP 1-selection mechanism that gives a 2-approximation.

In addition, 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 for 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

# References

[1] N. Alon, M. Feldman, A. D. Procaccia, and M. Tennenholtz. Strategyproof approximation mechanisms for location on networks. Manuscript, 2009.

[2] N. Alon and M. Naor. Coin-flipping games immune against linear-sized coalitions. *SIAM Journal on Computing*, 22:403–417, 1993.

[3] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley, 2nd edition, 2000.

[4] 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 (AAMAS)*, pages 1549–1552, 2008.

[5] S. Antonakopoulos. Fast leader-election protocols with bounded cheaters' edge. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 187–196, 2006.

[6] S. J. Brams and P. C. Fishburn. *Approval Voting*. Springer, 2nd edition, 2007.

[7] A. Cheng and E. Friedman. Sybilproof reputation mechanisms. In *Proceedings of the 3rd Workshop on the Economics of Peer-to-Peer Systems (P2PECON)*, pages 128–132, 2005.

[8] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[9] J. Cooper and N. Linial. Fast perfect-information leader-election protocols with linear immunity. *Combinatorica*, 15:319–332, 1995.

[10] O. Dekel, F. Fischer, and A. D. Procaccia. Incentive compatible regression learning. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 277–286, 2008.

[11] 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 (STOC)*, pages 644–652, 2006.

[12] U. Feige. Noncryptographic selection protocols. In *Proceedings of the 40th Symposium on Foundations of Computer Science (FOCS)*, pages 142–152, 1999.

[13] 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 (WINE)*, pages 878–886, 2005.

[14] 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 (STOC)*, pages 72–81, 2002.

[15] E. Friedman, P. Resnick, and R. Sami. Manipulation-resistant reputation systems. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 27. Cambridge University Press, 2007.

[16] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.

[17] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45:665–681, 1977.

[18] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55:242–269, 2006.

[19] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[20] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9. Cambridge University Press, 2007.

[21] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1–2):166–196, 2001.

[22] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce (ACM-EC)*, pages 177–186, 2009.

[23] 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.

[24] J. Schummer and R. V. Vohra. Mechanism design without money. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 10. Cambridge University Press, 2007.

[25] W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

# A    Proof of Theorem 4.1

For the first part of the theorem, consider an optimal set of $k$ agents (which might not be unique), and denote it by $K^* \subseteq N$. Let OPT be the sum of the indegrees of the agents in $K^*$, that is,

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

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

Consider some partition $\pi$ 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 $K_1^*$, that is,

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

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 $\lceil k/2 \rceil$ 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

$$\mathbb{E}\,[\,\text{MECH} \mid \pi\,] = \mathbb{E}\,[\,\text{MECH} \mid \pi \wedge T = \{1\}\,] \cdot \frac{1}{2} + \mathbb{E}\,[\,\text{MECH} \mid \pi \wedge T = \{2\}\,] \cdot \frac{1}{2}$$

$$\geq \left(\frac{\lceil k/2 \rceil}{k} \cdot d_1 + d_2\right) \cdot \frac{1}{2} + \left(\frac{\lfloor k/2 \rfloor}{k} \cdot d_1 + d_2\right) \cdot \frac{1}{2} \tag{3}$$

$$= \frac{d_1}{2} + d_2 \geq \frac{d_1 + d_2}{2}.$$

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 OPT/2. Formally, for a partition $\pi$, let $S_1^\pi$ and $S_2^\pi$ 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

$$\sum_\pi \Pr[\,\pi\,] \cdot d^\pi = \frac{\text{OPT}}{2}. \tag{4}$$

We can now conclude that

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

Figure 3: An illustration of the proof of Theorem 4.1, for $n = 8$ and $k = 4$. In the given graph $G$, 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 $d_1 = d_2 = 1$.

where the second transition follows from (3) and the third transition follows from (4).

We now turn to the second part of the theorem. For ease of exposition, we will omit the various floors and ceilings from the proof, as we are looking for an asymptotic result. We employ one additional idea: 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 $K^* = \{1, \ldots, k\}$. Given $i \in K^*$, we define a random variable $Z_i$ that depends on the random partition of $N$ to $S_1, \ldots, S_{k^{1/3}}$ as follows:

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

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

$$\mathbb{E}[\text{MECH}] = \sum_{s_1, \ldots, s_k} \mathbb{E}[\text{MECH} \mid Z_1 = s_1, \ldots, Z_k = s_k] \cdot \Pr[Z_1 = s_1, \ldots, Z_k = s_k], \qquad (5)$$

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 (3), we get

$$\mathbb{E}[\text{MECH} \mid Z_1 = s_1, \ldots, Z_k = s_k] \geq \mathbb{E}\left[\sum_{i \in K^*} d_i' \sigma_{s_i} \mid Z_1 = s_1, \ldots, Z_k = s_k\right]$$
$$= \sum_{i \in K^*} \mathbb{E}[d_i' \sigma_{s_i} \mid Z_1 = s_1, \ldots, Z_k = s_k]. \qquad (6)$$

14

We wish to obtain an explicit expression for $\mathbb{E}\,[\,d_i'\sigma_{s_i} \mid Z_1 = s_1, \ldots, Z_k = s_k\,]$. For $i \in N$ and $S \subseteq N$, let

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

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

$$\mathbb{E}\,[\,d_i'\sigma_{s_i} \mid Z_1 = s_1, \ldots, Z_k = s_k\,] = \left(\frac{k - 1 - s_i}{k - 1} \cdot \deg(i, K^*) + \frac{k^{1/3} - 1}{k^{1/3}} \cdot \deg(i, N \setminus K^*)\right) \cdot \sigma_{s_i}. \quad (7)$$

Indeed, 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 (7) is independent of $s_j$ for all $j \neq i$.

Combining (5), (6), and (7), and reversing the order of summation, we conclude that

$$\mathbb{E}\,[\,\text{MECH}\,] \geq \sum_{i \in K^*} \sum_{s_1, \ldots, s_k} \Pr[\,Z_1 = s_1, \ldots, Z_k = s_k\,] \cdot$$

$$\left(\frac{k - 1 - s_i}{k - 1} \cdot \deg(i, K^*) + \frac{k^{1/3} - 1}{k^{1/3}} \cdot \deg(i, N \setminus K^*)\right) \cdot \sigma_{s_i}$$

$$= \sum_{i \in K^*} \sum_{s=0}^{k-1} \Pr[\,Z_i = s\,] \cdot \left(\frac{k - 1 - s}{k - 1} \cdot \deg(i, K^*) + \frac{k^{1/3} - 1}{k^{1/3}} \cdot \deg(i, N \setminus K^*)\right) \cdot \sigma_s$$

$$= \sum_{i \in K^*} \sum_{s=0}^{k-1} \Pr[\,Z_i = s\,] \cdot \frac{k - 1 - s}{k - 1} \cdot \deg(i, K^*) \cdot \sigma_s +$$

$$\sum_{i \in K^*} \sum_{s=0}^{k-1} \Pr[\,Z_i = s\,] \cdot \frac{k^{1/3} - 1}{k^{1/3}} \cdot \deg(i, N \setminus K^*) \cdot \sigma_s.$$

On the other hand, we have that

$$\text{OPT} = \sum_{i \in K^*} \left(\deg(i, K^*) + \deg(i, N \setminus K^*)\right) = \sum_{i \in K^*} \deg(i, K^*) + \sum_{i \in K^*} \deg(i, N \setminus K^*).$$

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 conclude that

$$\frac{\text{OPT}}{\mathbb{E}\,[\,\text{MECH}\,]} \leq \frac{\sum_{i \in K^*} \deg(i, K^*) + \sum_{i \in K^*} \deg(i, N \setminus K^*)}{\sum_{i \in K^*} \left(1 - \mathcal{O}\left(\frac{1}{k^{1/3}}\right)\right) \deg(i, K^*) + \sum_{i \in K^*} \left(1 - \mathcal{O}\left(\frac{1}{k^{1/3}}\right)\right) \deg(i, N \setminus K^*)}$$

$$= \frac{1}{\left(1 - \mathcal{O}\left(\frac{1}{k^{1/3}}\right)\right)} = 1 + \mathcal{O}\left(\frac{1}{k^{1/3}}\right).$$

15

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,

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

In order to bound the probabilities on the right hand side of (10) we employ the following version of the Chernoff bounds (see, e.g., [3], Theorem A.1.11).

**Lemma A.1.** *Let $X_1, ..., X_k$ be i.i.d. Bernoulli trials, $\Pr[\, X_i = 1 \,] = p$ for $i = 1, \ldots, k$, and denote $X = \sum_{i=1}^{k} X_i$. 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 if we are interested in an upper bound on the probability of the sum being greater than some given value. Using Lemma A.1 with $\lambda = xk^{1/3}$ and $p = 1/k^{1/3}$ we get

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

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

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

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

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

and as before,

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

Finally,

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

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

We can thus bound this sum from above as before using (11). This completes the proof of Theorem 4.1.

# B  The Edge Scan Mechanism

In Theorem 3.1 we have seen that a deterministic SP $k$-selection mechanism cannot give a bounded approximation ratio. We now show that if we are allowed to choose *at most* $k$ agents, for $k \ge 2$, it is possible to design an SP mechanism with a bounded approximation ratio. As noted in Section 3, it is sufficient to select a subset with an incoming edge, if one exists.

Intuitively, the mechanism, which we refer to as the *Edge Scan Mechanism*, first orders the agents from left to right according to their lexicographic ordering. The mechanism 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. Similarly, the mechanism scans the agents from right to left until it finds an edge that is directed to the left, and also selects the agent that this edge is pointing at. An example is shown in Figure 4. What follows is a more formal specification of the mechanism.

**The Edge Scan Mechanism.**

1. Partition $E$ into $E_1 = \{(i,j) \in E : i < j\}$ and $E_2 = \{(i,j) \in E : i > j\}$.

2. If $E_1 \ne \emptyset$, let $i \in N$ be the minimum index such that there exists $j \in N$ with $(i,j) \in E_1$; add to the subset the minimum $j$ such that $(i,j) \in E_1$. Otherwise, add agent $n$ to the subset.

3. If $E_2 \ne \emptyset$, let $i \in N$ be the maximum index such that there exists $j \in N$ with $(i,j) \in E_2$; add to the subset the maximum $j$ such that $(i,j) \in E_2$. Otherwise, add agent 1 to the subset.

The Edge Scan Mechanism is clearly SP. Indeed, agent $i$ cannot benefit from adding outgoing edges, since these edges would only point at some other agent. It also cannot benefit from removing outgoing edges. Informally, if the mechanism reaches the point in the scan (from left to right or right to left) where the agent's vote is taken into account, then it is too late for agent $i$ itself to be elected.

Figure 4: Example for the Edge Scan Mechanism. Given this graph, the mechanism would select agent 4 in the scan from left to right, and agent 3 in the scan from right to left, so the subset of agents selected by the mechanism is $\{3, 4\}$.

Moreover, either $E_1$ or $E_2$ will contain an edge given that there is at least one edge in the graph, and the Edge Scan Mechanism is guaranteed to select an agent with an incoming edge in this case. It therefore achieves a finite approximation ratio, although this ratio can be as bad as $\Omega(nk)$.

Crucially, the agents selected in both steps of the mechanism can be one and the same; in this case the mechanism would return a singleton subset. A curious implication of Theorem 3.1 is that such a selection cannot be completed deterministically and in a strategyproof way to obtain a subset of size two.

# C  The Sliding Partition Mechanism

We first give a formal specification of the mechanism, informally presented in Section 4.

**The Sliding Partition Mechanism.**

1. Let $S = \emptyset$.

2. While $|S| < n - 1$, choose $i \notin S$ that has minimum indegree based on edges from agents in $S$, breaking ties randomly. Let $S = S \cup \{i\}$.

3. Select the agent in $N \setminus \{S\}$.

When an agent is added to $S$, we say that it is *eliminated*. It is easy to see that this mechanism is SP. Indeed, only the outgoing edges of eliminated agents are taken into account at any stage. Once an agent is eliminated, it no longer has a chance to be selected, therefore it is indifferent to the outcome of the mechanism.

Another interesting observation is that the Sliding Partition Mechanism gives a 2-approximation for the example where the analysis of the 2-RP mechanism is tight: a graph with only one edge. Indeed, if $G$ has one edge $(i, j)$, then $j$ is certainly elected once $i$ is eliminated (since then it is the only agent in $N \setminus S$ with an incoming edge from $S$), and $i$ is eliminated before $j$ with probability $1/2$.

Unfortunately, it is possible to construct a graph where the mechanism does very poorly. For this, consider a tree with agent 1 at the root. There is a set $T \subset N$ of size $n^{3/5}$ of agents with outgoing edges to 1, that is, $\deg(1) = n^{3/5}$. In addition, each agent in $T$ has $n^{2/5}$ incoming edges from agents in $N \setminus (\{1\} \cup T)$. The agents in $N \setminus (\{1\} \cup T)$ have an indegree of zero.

Notice that while there are agents in $N \setminus S$ that have no incoming edges from $S$, the mechanism selects one of these agents uniformly at random and eliminates it. Consider the first point in time $t_0$ when all the agents in $T \setminus S$ that were not yet eliminated have at least one incoming edge from $S$; we can assume without loss of generality that at this point agent 1 has not been eliminated.

We claim that if less than $n^{2/5}$ agents from $T$ have been eliminated at time $t_0$, then agent 1 is guaranteed to be eliminated later on. Indeed, starting at $t_0$, the remaining agents in $N \setminus (\{1\} \cup T)$ are eliminated one after the other (in some random order), because all of them have indegree zero from $S$, while all other remaining agents have at least one incoming edge from $S$. After all the agents in $N \setminus (\{1\} \cup T)$ have been eliminated, each agent in $T \setminus S$ has $n^{2/5}$ incoming edges from $S$. By assumption, agent 1 has less, and is eliminated next.

We now claim that with high probability, agent 1 has less than $n^{2/5}$ incoming edges from $S$ at time $t_0$. Each agent $i \in T$ contributes an edge to 1 at time $t_0$ if and only if it is eliminated before any of the agents in its incoming neighborhood; this happens with probability roughly $1/n^{2/5}$. Therefore, by linearity of expectation, the expected number of edges from $S$ to agent 1 at time $t_0$ is roughly $n^{1/5}$. The claim now follows directly from Chernoff's inequality.

We conclude that the approximation ratio provided by the Sliding Partition Mechanism cannot be smaller than $\Omega(n^{1/5})$. By optimizing the parameters of the example, it is possible to obtain an even stronger lower bound.