Impartial Selection and the Power of Up to Two Choices

ANTJE BJELDE, Humboldt-Universität zu Berlin, Germany
FELIX FISCHER, University of Glasgow, UK
MAX KLIMM, Humboldt-Universität zu Berlin, Germany

We study mechanisms that select members of a set of agents based on nominations by other members and that are impartial in the sense that agents cannot influence their own chance of selection. Prior work has shown that deterministic mechanisms for selecting any fixed number k of agents are severely limited and cannot extract a constant fraction of the nominations of the k most highly nominated agents. We prove here that this impossibility result can be circumvented by allowing the mechanism to sometimes but not always select fewer than k agents. This added flexibility also improves the performance of randomized mechanisms, for which we show a separation between mechanisms that make exactly two or up to two choices and give upper and lower bounds for mechanisms allowed more than two choices.

CCS Concepts: • **Theory of computation → Algorithmic game theory and mechanism design**;

Additional Key Words and Phrases: Mechanism design, impartial selection

ACM Reference format:
Antje Bjelde, Felix Fischer, and Max Klimm. 2017. Impartial Selection and the Power of Up to Two Choices. *ACM Transactions on Economics and Computation* 0, 0, Article 1 (April 2017), 20 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

We consider the setting of impartial selection first studied by Alon et al. [1] and by Holzman and Moulin [8]. The goal in this setting is to select members of a set of agents based on nominations cast by other members of the set, under the assumption that agents will reveal their true opinion about other agents as long as they cannot influence their own chance of selection. The assumption of impartiality seems justified, and is routinely made, in many situations where a strong correlation exists between expertise and self-interest, like the selection of representatives from within a group and the use of peer review in the allocation of funding and scientific or academic credit.

Formally, the impartial selection problem can be modeled by a directed graph with n vertices, one for each agent, in which edges correspond to nominations. A selection mechanism then chooses, possibly using randomization, a set of vertices for any given graph, and it is impartial if the chances of a particular vertex to be chosen do not depend on its outgoing edges. As impartiality may prevent us from simply selecting the vertices with maximum indegree, corresponding to the most highly

The research was carried out in the framework of MATHEON supported by the Einstein Foundation Berlin. Valuable feedback from the anonymous referees is gratefully acknowledged.
Author's addresses: A. Bjelde and M. Klimm, Wirtschaftswissenschaftliche Fakultät, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany. Email: {antje.bjelde,max.klimm}@hu-berlin.de. F. Fischer, School of Computing Science, University of Glasgow, Sir Alwyn Williams Building, Glasgow G12 8QQ, UK. Email: fischerf@dcs.gla.ac.uk.
 XXXX-XXXX/2017/4-ART1 $15.00
DOI: 10.1145/nnnnnnn.nnnnnnn

k	deterministic	randomized exact	randomized
1	0	$\frac{1}{2}$	$\frac{1}{2}$
2	$\frac{1}{2}$	$\left[\frac{7}{12}, \frac{2}{3}\right]$	$\left[\frac{2}{3}, \frac{3}{4}\right]$
$\vdots$			
k	$\left[\frac{1}{k}, \frac{k-1}{k}\right]$	$\left[\frac{k}{k+1}\left(1-(\frac{k-1}{k})^{k+1}\right), \frac{k+1}{k+2}\right]$	$\left[\frac{k}{k+1}\left(1-(\frac{k-1}{k})^{k+1}\right), \frac{k+1}{k+2}\right]$
$\vdots$			
$n-2$	$\left[\frac{1}{k}, \frac{k-1}{k}\right]$	$\left[\frac{k}{k+1}\left(1-(\frac{k-1}{k})^{k+1}\right), \frac{7k^3+5k^2-6k+12}{7k^3+13k^2-2k}\right]$	$\left[\frac{k}{k+1}\left(1-(\frac{k-1}{k})^{k+1}\right), \frac{k+1}{k+2}\right]$
$n-1$	$\left[\frac{1}{k}, \frac{k-1}{k}\right]$	$\left[\frac{k}{k+1}\left(1-(\frac{k-1}{k})^{k+1}\right), \frac{k}{k+1}\right]$	$\left[\frac{k}{k+1}\left(1-(\frac{k-1}{k})^{k+1}\right), \frac{2k}{2k+1}\right]$

Table 1. Bounds on α for α-optimal impartial selection of at most or exactly k agents. Deterministic exact mechanisms, not shown in the table, cannot be α-optimal for any $\alpha > 0$.

nominated agents, it is natural to instead approximate this objective. For an integer k, a selection mechanism is called a k-selection mechanism if it selects at most k vertices of any input graph. We call a k-selection mechanism *exact* if it always selects exactly k agents. A k-selection mechanism is called *α-optimal*, for $\alpha \leq 1$, if for any input graph the sum of indegrees of the selected vertices is at least α times the sum of the k largest indegrees.

In prior work, a striking separation was shown between mechanisms that do not use randomness and those that do. On the one hand, no deterministic exact α-optimal mechanism exists for selecting any fixed number of agents and any $\alpha > 0$ [1]. On the other, a mechanism that aligns the agents along a random permutation from left to right and selects a single agent with a maximum number of nominations from its left achieves a bound of $\alpha = 1/2$ [7]. This bound is in fact best possible subject to impartiality [1].

Our Contribution. We show here that a relaxation of exactness is another remedy to the strong impossibility result concerning exact deterministic k-selection mechanisms in that it enables the design of α-optimal mechanisms for constant α. Specifically, for $k = 2$, running the permutation mechanism on a fixed instead of a random permutation but selecting an agent for each direction of that permutation is $1/2$-optimal. The factor of $1/2$ is again best possible. Flexibility in the exact number of selected agents is beneficial also in the realm of randomized impartial mechanisms: given a set of three agents, for example, a $3/4$-optimal mechanism exists selecting two agents or fewer, whereas the best mechanism selecting exactly two agents is only $2/3$-optimal. For 2-selection from an arbitrary number of agents, we give a randomized exact $7/12$-optimal mechanism and a randomized $2/3$-optimal mechanism that is not exact. Finally we provide upper and lower bounds on the performance of mechanisms allowed to make more than two choices. A summary of our current state of knowledge is shown in Table 1.

Related Work. The theory of impartial decision making was first considered by de Clippel et al. [6], for the case of a divisible resource to be shared among a set of agents. The difference between divisible and indivisible resources disappears for randomized mechanisms, but the mechanisms of de Clippel et al. allow for fractional nominations and do not have any obvious consequences for our setting. Impartial selection is a rather fundamental problem in social choice theory, with applications ranging from the selection of committees to academic peer review. The problem we consider here was first studied by Holzman and Moulin [8] and Alon et al. [1], the articles of

Holzman and Moulin and of Fischer and Klimm [7] provide a good introduction to its history and early literature. When agents are interested purely in their own selection the problem can be viewed as an example of mechanism design without money, an agenda put forward by Procaccia and Tennenholtz [13]. In peer review the need for impartiality is only one of a number of issues along with information elicitation and incentivization of effort, and a natural approach would be to combine our mechanisms with mechanisms seeking to achieve the other goals [e.g., 17, 18]. Other authors have taken a more holistic view of peer review and peer selection and have aimed for more practical and more heuristic mechanisms [2, 9]. Tamura and Ohseto [16] were the first to consider impartial mechanisms selecting more than one agent and showed that these can circumvent some of the impossibility results of Holzman and Moulin. An axiomatic characterization of the mechanisms was later given by Tamura [15]. Mackenzie [10] gave a characterization of symmetric randomized selection mechanisms for the special case where each agent nominates exactly one other agent. Inspiration for our title, and indeed for relaxing the requirement to always select the same number of agents, comes from the power of multiple choices in load balancing, where even two choices can lead to dramatically lower average load [e.g., 12]. The related concept of resource augmentation, first used by Sleator and Tarjan [14], is a common technique in the analysis of online algorithms and has also been applied to a problem in mechanism design [5]. Mackenzie [11] recently studied the relationship between impartiality, exactness, and randomization for various mechanisms used over the centuries in electing the pope.

Open Problems. With the exception of mechanisms that are asymptotically optimal, when many agents are selected [1] or when agents receive many nominations [4], only very little was previously known about the impartial selection of more than one agent. Our understanding of 2-selection is now much better, with some room for improvement in the case of randomized mechanisms. About k-selection for $k > 2$, in particular about deterministic mechanisms for this task, we still know relatively little. This lack of understanding is witnessed by the fact that the optimal deterministic mechanism selecting up to two agents, one for each direction of a permutation, does not generalize in any obvious way to the selection of more than two agents. We may also hope for stronger techniques to bound the power of randomized mechanism that are universally impartial in the sense that they can be obtained as a convex combination of deterministic impartial mechanisms, and to design mechanisms that are impartial but not universally impartial. We will see in Section 4 that optimal mechanisms in the latter category may exhibit rather unintuitive nonmonotonicity properties, which in turn complicates their analysis. Meanwhile, the existence of a near-optimal mechanism in the limit of many selected agents suggests that the upper rather than the lower bounds may be correct. Table 1 illustrates that the relaxation of exactness can benefit both deterministic and randomized mechanisms, and that randomization can be beneficial independently of exactness. It is not clear, however, whether either of these statements is true for all values of n and k.

2 PRELIMINARIES

For $n \in \mathbb{N}$, let

$$\mathcal{G}_n = \left\{(N, E) : N = \{1, \dots, n\}, E \subseteq \{(i, j) \in N \times N : i \neq j\}\right\}$$

be the set of directed graphs with n vertices and no loops. Let $\mathcal{G} = \bigcup_{n \in \mathbb{N}} \mathcal{G}_n$. For $G = (N, E) \in \mathcal{G}$ and $S, X \subseteq N$ let

$$\delta_S^-(X, G) = |\{(j, i) \in E : G = (N, E), j \in S, i \in X\}|$$

denote the sum of indegrees of vertices in X from vertices in S. We use $\delta^-(X, G)$ as a shorthand for $\delta_N^-(X, G)$ and denote by $\Delta_k(G) = \max_{X \subseteq N, |X|=k} \delta^-(X, G)$. When $X = \{i\}$ for a single vertex i, we

write $\delta_S^-(i,G)$ instead of $\delta_S^-(\{i\},G)$. Most of the time, the graph G will be clear from context. We then write $\delta_S^-(X)$ instead of $\delta_S^-(X,G)$, $\delta^-(X)$ instead of $\delta^-(X,G)$, and Δ_k instead of $\Delta_k(G)$.

For $n,k \in N$, let $\mathcal{X}_n = \{X : X \subseteq \{1,\dots,n\}\}$ be the set of subsets of the first n natural numbers and let $\mathcal{X}_{n,k} = \{X \in \mathcal{X}_n : |X| = k\}$ be the subset of these sets with cardinality k. A *k-selection mechanism* for $\mathcal{G}$ is then given by a family of functions $f : \mathcal{G}_n \to [0,1]^{\bigcup_{l=0}^{k} \mathcal{X}_{n,l}}$ that maps each graph to a probability distribution on subsets of at most k of its vertices. In a slight abuse of notation, we use f to refer to both the mechanism and individual functions from the family.

We call mechanism f *deterministic* if $f(G) \in \{0,1\}^{\bigcup_{l=0}^{k} \mathcal{X}_{n,l}}$, i.e., if $f(G)$ puts all probability mass on a single set for all $G \in \mathcal{G}$; and *exact* if $(f(G))_X = 0$ for every $n \in \mathbb{N}$, $G \in \mathcal{G}_n$, and $X \in \mathcal{X}_n$ with $|X| < k$, i.e., if the mechanism never selects a set X of vertices with strictly less than k vertices.

Mechanism f is *impartial* on $\mathcal{G}' \subseteq \mathcal{G}$ if on this set of graphs the probability of selecting vertex i does not depend on its outgoing edges, i.e., if for every pair of graphs $G = (N,E)$ and $G' = (N,E')$ in $\mathcal{G}'$ and every $i \in N$, $\sum_{X \in \mathcal{X}_n, i \in X} (f(G))_X = \sum_{X \in \mathcal{X}_n, i \in X} (f(G'))_X$ whenever $E \setminus (\{i\} \times N) = E' \setminus (\{i\} \times N)$. Note that while impartiality requires the outgoing edges of a vertex i to have no influence at all on whether i is selected or not, they may influence both the number and the identities of other vertices selected. All mechanisms we consider are impartial on $\mathcal{G}$, and we simply refer to such mechanisms as impartial mechanisms.

Finally, a k-selection mechanism f is *α-optimal* on $\mathcal{G}' \subseteq \mathcal{G}$, for $\alpha \leq 1$, if for every graph in $\mathcal{G}'$ the expected sum of indegrees of the vertices selected by f differs from the maximum sum of indegrees for any k-subset of the vertices by a factor of at most α, i.e., if

$$\inf_{\substack{G \in \mathcal{G}' \\ \Delta_k(G) > 0}} \frac{\mathbb{E}_{X \sim f(G)}[\delta^-(X,G)]}{\Delta_k(G)} \geq \alpha.$$

We call a mechanism α-optimal if it is α-optimal on $\mathcal{G}$.

For randomized mechanisms, and as far as impartiality and α-optimality are concerned, we can restrict attention to mechanisms that are *symmetric*, i.e., invariant with respect to renaming of the vertices [e.g., 7]. It may further be convenient to view a k-selection mechanism as assigning probabilities to vertices rather than sets of vertices, with the former summing to at most k or exactly k for each graph. By the Birkhoff-von Neumann theorem [3], the two views are equivalent in the following way.

Lemma 2.1. *Let $n \in \mathbb{N}$, $p \in [0,1]^n$, and $m = \sum_{i=1}^n p_i$. Then there exists a random variable Y with values in $[0,1]^{\mathcal{X}_{n,\lfloor m \rfloor} \cup \mathcal{X}_{n,\lceil m \rceil}}$ such that for all $i \in \{1,\dots,n\}$, $\sum_{X \in \mathcal{X}_n, i \in X} \mathbb{P}[Y = X] = p_i$.*

Proof. First consider the case where m is an integer, and let $M = \{1,\dots,m\}$ and $\bar{M} = \{m+1,\dots,n\}$. Since $\sum_{i=1}^n p_i = m$, there exists $Q \in [0,1]^{n \times m}$ such that for all $i \in \{1,\dots,n\}$, $\sum_{j \in M \cup \bar{M}} q_{ij} = 1$ and $\sum_{j \in M} q_{ij} = p_i$, and for all $j \in M \cup \bar{M}$, $\sum_{i=1}^n q_{ij} = 1$. Thus Q is doubly stochastic, and by the Birkhoff-von Neumann theorem can be written as a convex combination of permutation matrices. For each individual permutation matrix R there then exists a set $X \in \mathcal{X}_{n,m}$ such that $r_{ij} = 1$ for some $j \in M$ if and only if $i \in X$, which shows the claim.

When m is an arbitrary number, we can write p as a convex combination of two vectors p' and p'' such that $\sum_{i=1}^n p'_i = \lfloor m \rfloor$ and $\sum_{i=1}^n p''_i = \lceil m \rceil$. The claim then follows by applying the above reasoning independently to p' and p''. □

3 DETERMINISTIC MECHANISMS

Focusing on the exact case, Alon et al. showed that deterministic k-selection mechanisms cannot be α-optimal for any $k \in \{1,\dots,n-1\}$ and any $\alpha > 0$. This result is a rather simple observation for $k = 1$, but quite surprising when $k > 1$. For $(n-1)$-selection in particular, any deterministic

ALGORITHM 1: The bidirectional permutation mechanism, using extraction mechanism Ξ_π

Input: Graph $G = (N, E)$
Output: Set $\{i_1, i_2\} \subseteq N$ of at most two vertices
Let $\pi = (1, \ldots, n)$;
$i_1 := \Xi_\pi(G)$; ▷ select vertex based on forward edges
$i_2 := \Xi_{\bar{\pi}}(G)$; ▷ select vertex based on backward edges
return $\{i_1, i_2\}$;

mechanism that is both exact and impartial must sometimes exclude precisely the unique vertex with positive indegree. While it is not difficult to convince ourselves that a relaxation of exactness is not helpful in the case of 1-selection, we will exhibit momentarily a deterministic impartial mechanism that for any graph selects either one or two vertices whose overall indegree is at least the largest indegree of any vertex in the graph.

There are two ways to interpret this result. Since the largest indegree is at least half of the sum of the two largest indegrees, relaxing exactness allows us to circumvent the strong lower bound of Alon et al. when $k = 2$. Alternatively, for $k = 1$, the tradeoff between impartiality and quality of the outcome disappears if one is allowed to sometimes but not always select an additional vertex. This kind of resource augmentation result, comparing an optimal algorithm to one from a restricted class that is given additional resources, is commonly used in the analysis of online algorithms and has recently also been applied to truthful mechanisms for facility assignment [5].

To explain our mechanism in detail we need some additional notation. Let $N = \{1, \ldots, n\}$. For a graph $G = (N, E)$ and a permutation $\pi = (\pi_1, \ldots, \pi_n)$ of N, denote by

$$E_\pi = \left\{(u, v) \in E : \pi_i = u, \pi_j = v \text{ for some } i, j \text{ with } 1 \leq i < j \leq n\right\}$$

the set of *forward edges* of G with respect to π. Denote by $\bar{\pi}$ the permutation obtained by reading π backwards, such that $\bar{\pi}_i = \pi_{n+1-i}$ for $i = 1, \ldots, n$. Finally, for a permutation π and $j \in \{1, \ldots, n\}$, let $\pi_{<j} = \{\pi_1, \pi_2, \ldots, j\} \setminus \{j\}$ denote the set of vertices in the prefix of π up to but not including j.

The first mechanism we consider, which we call the bidirectional permutation mechanism, considers the vertices one by one according to a fixed permutation π and in each step compares the current vertex π_j to a single candidate vertex π_l with $l < j$. In determining the indegree of the candidate vertex π_l it takes into account the outgoing edges of vertices $\pi_1, \ldots, \pi_{l-1}$. For the indegree of the current vertex π_j it takes into account the outgoing edges of vertices $\pi_1, \ldots, \pi_{j-1}$, with the exception of π_l. If the latter is greater than or equal to the former, π_j becomes the new candidate, and the candidate after the final step is the first vertex selected by the mechanism. The same procedure is then applied with permutation $\bar{\pi}$ to find a second vertex. A formal description of the bidirectional permutation mechanism is given as Algorithm 1. It is formulated in terms of Algorithm 2, which we call the extraction mechanism and which is identical to a mechanism of Fischer and Klimm except for its use of a given permutation rather than a random one.

It is worth noting that the bidirectional permutation mechanism may select only one vertex, namely if the same vertex is chosen for both directions of the permutation. This happens for example in the graph of Figure 1(a).

To see that the bidirectional permutation mechanism is impartial, we first note that this is true for a single run of the extraction mechanism. Indeed, the outcome of the latter is influenced by the outgoing edges of any given vertex only when that vertex can no longer be selected.

Lemma 3.1. *The extraction mechanism is impartial.*

ALGORITHM 2: The extraction mechanism Ξ_π

Input: Graph $G = (N, E)$, permutation $(\pi_1, \ldots, \pi_n)$ of N
Output: Vertex $i \in N$
Set $i := \pi_1$, $d := 0$; ▷ candidate vertex and its indegree from its left
for $j = 2, \ldots, n$ **do**
  if $\delta^-_{\pi_{<\pi_j}\setminus\{i\}}(\pi_j) \geq d$ **then** ▷ compare current vertex and candidate
    Set $i := \pi_j$, $d := \delta^-_{\pi_{<\pi_j}}(\pi_j)$; ▷ current vertex becomes new candidate
  end
end
return i;

(a) (b)

Fig. 1. Graphs for which the bidirectional permutation mechanism returns only one vertex (a) and is only 1/2-optimal (b).

Impartiality of the bidirectional permutation mechanism then follows because the union of the results of k impartial 1-selection mechanisms yields an impartial k-selection mechanism.

Lemma 3.2. *Let $f_1, \ldots, f_k$ be impartial 1-selection mechanisms. Then the mechanism that selects the vertices selected by at least one of the mechanisms $f_1, \ldots, f_k$ is an impartial k-selection mechanism.*

Proof. By impartiality of f_l, for $l = 1, \ldots, k$, the outgoing edges of a vertex do not influence whether this vertex is selected by f_l. This holds for any l and any vertex, so it also holds for the mechanism that selects the vertices selected by at least one of the mechanisms. □

We now proceed to show that the bidirectional permutation mechanism is 1/2-optimal, starting from the observation that the vertex selected by Ξ_π has a maximum number of incoming forward edges with respect to π.

Lemma 3.3. *If $i = \Xi_\pi(G)$, then $\delta^-_{\pi_{<i}}(i, G) = \max_{j=1,\ldots,n}\{\delta^-_{\pi_{<j}}(j, G)\}$.*

Proof. Let $d^* = \max_{j=1,\ldots,n}\{\delta^-_{\pi_{<j}}(j)\}$, and let i^* be an arbitrary vertex with $\delta^-_{\pi_{<i^*}}(i^*) = d^*$. When i^* is considered by the mechanism, so are at least $d^* - 1$ of its incoming forward edges, one of which may originate from the current candidate i, i.e., $\delta^-_{\pi_{<i^*}\setminus\{i\}}(i^*) \in \{d^* - 1, d^*\}$.

If $\delta^-_{\pi_{<i^*}\setminus\{i\}}(i^*) = d^*$ or both $\delta^-_{\pi_{<i^*}\setminus\{i\}}(i^*) = d^* - 1$ and $\delta^-_{\pi_{<i}}(i) \leq d^* - 1$, then i^* becomes the new candidate. Since d will be set to d^*, any other vertex that possibly becomes a candidate after i^* has d^* incoming forward edges as well, establishing the claim for this case.

If, on the other hand, both $\delta^-_{\pi_{<i^*}\setminus\{i\}}(i^*) = d^* - 1$ and $\delta^-_{\pi_{<i}}(i) = d^*$, then i remains the candidate. As $d = d^*$ any further candidate has d^* incoming forward edges as in the first case. □

Theorem 3.4. *The bidirectional permutation mechanism is impartial and 1/2-optimal.*

Proof. Impartiality follows directly from Lemma 3.1 and Lemma 3.2.

Now consider a graph $G = (N, E)$, a vertex i^* with $\delta^-(i^*) = \Delta_1$, and let $i_1 = \Xi_\pi(G)$ and $i_2 = \Xi_{\bar\pi}(G)$. By Lemma 3.3, $\delta^-_{\pi_{<i_1}}(i_1) \geq \delta^-_{\pi_{<i^*}}(i^*)$ and $\delta^-_{\bar\pi_{<i_2}}(i_2) \geq \delta^-_{\bar\pi_{<i^*}}(i^*)$, regardless of whether

ALGORITHM 3: The 2-partition mechanism with permutation

Input: Graph $G = (N, E)$ with $n \geq 2$
Output: Vertices $i_1, i_2 \in N$.
Assign each $i \in N$ to A_1 or A_2 independently and uniformly at random;
Choose a permutation $(\pi_1, \ldots, \pi_n)$ of N uniformly at random;
for $j = 1, 2$ **do**
  $i_j := \Xi_{\pi, A_j}(G)$; ▷ select one vertex from each of the two sets
end
▷ if one set is empty, select 2nd vertex from other set
if $A_2 = \emptyset$ **then** choose i_2 uniformly at random from $A_1 \setminus \{i_1\}$;
if $A_1 = \emptyset$ **then** choose i_1 uniformly at random from $A_2 \setminus \{i_2\}$;
return $\{i_1, i_2\}$;

ALGORITHM 4: The extraction mechanism $\Xi_{\pi, A}$ restricted to a set $A \subseteq N$

Input: Graph $G = (N, E)$, permutation $(\pi_1, \ldots, \pi_n)$ of N, set $A \subseteq N$
Output: Vertex $i \in N$
Set $i := \pi_1$, $d := 0$; ▷ candidate vertex and indegree from its left
for $j = 2, \ldots, n$ **do**
  $S := (N \setminus A) \cup (\pi_{<\pi_j} \setminus \{i\})$; ▷ vertices that cannot be selected
  if $\pi_j \in A$ **and** $\delta^-_S(\pi_j) \geq d$ **then** ▷ compare current vertex and candidate
    Set $i := \pi_j$, $d := \delta^-_{S \cup \{i\}}(\pi_j)$; ▷ current vertex becomes new candidate
  end
end
return i;

$i_1 \neq i_2$ or $i_1 = i_2$. Thus

$$\begin{aligned}\delta^-(\{i_1, i_2\}) &\geq \delta^-_{\pi_{<i_1}}(i_1) + \delta^-_{\bar{\pi}_{<i_2}}(i_2) \\ &\geq \delta^-_{\pi_{<i^*}}(i^*) + \delta^-_{\bar{\pi}_{<i^*}}(i^*) = \delta^-(i^*) = \Delta_1 \geq \frac{1}{2}\Delta_2,\end{aligned}$$

as claimed. □

To see that the analysis is tight, consider the graph in Figure 1(b). For this graph, the mechanism selects vertices π_2 and π_1 with an overall indegree of 1, while the maximum overall indegree of a set of two vertices is 2. We will see later, in Theorem 6.1, that the bound of $1/2$ is in fact best possible.

4 RANDOMIZED MECHANISMS

In light of the results of the previous section, it is natural to ask whether a relaxation of exactness enables better bounds also for randomized mechanisms. We answer this question in the affirmative and give the first nontrivial bounds for both exact and inexact 2-selection mechanisms, as well as an example that shows a strict separation between the two classes.

We begin by considering an exact mechanism, which we call the 2-partition mechanism with permutation. The mechanism randomly partitions the set of vertices into two sets A_1 and A_2 such that $\mathbb{P}[i \in A_1] = \mathbb{P}[i \in A_2] = 1/2$ for all $i \in N$, $A_1 \cup A_2 = N$, and $A_1 \cap A_2 = \emptyset$. It then selects one vertex from each of the sets by applying the extraction mechanism with a random permutation, while also taking into account incoming edges from the respective other set. Algorithm 3 is a formal

description of the mechanism. It uses a restricted version of the extraction mechanism, given as Algorithm 4 and denoted $\Xi_{\pi,A}$ for a set $A \subseteq N$. The properties of the latter can be summarized in terms of the following two results.

Lemma 4.1. *The restricted extraction mechanism is impartial.*

Proof. When A is empty, the mechanism selects the same vertex for any graph and therefore is impartial. Otherwise the first vertex in A to appear in π becomes a candidate and only vertices from A are considered thereafter, so the mechanisms selects a vertex from A. Moreover, the mechanism only takes into account outgoing edges of vertices that can no longer be selected, either because they are not in A or because they have already been considered and are not currently the candidate. This directly implies impartiality. □

Lemma 4.2. *If* $i = \Xi_{\pi,A}(G)$, *then* $\delta^-_{(N\setminus A)\cup\pi_{<i}}(i,G) \geq \max_{j\in A}\{\delta^-_{(N\setminus A)\cup\pi_{<j}}(j,G)\}$.

Proof. The statement is trivial in case $A = \emptyset$. Otherwise, let $i^* \in \arg\max_{j\in A}\{\delta^-_{(N\setminus A)\cup\pi_{<j}}(j)\}$. Analogously to the proof of Lemma 3.3, we consider the iteration in which the mechanism decides whether i^* should become the new candidate. Let i be the candidate at the beginning of that iteration.

If $\delta^-_{(N\setminus A)\cup\pi_{<i^*}\setminus\{i\}}(i^*) = d^*$ or both $\delta^-_{(N\setminus A)\cup\pi_{<i^*}\setminus\{i\}}(i^*) = d^* - 1$ and $\delta^-_{(N\setminus A)\cup\pi_{<i}}(i) \leq d^* - 1$, then i^* becomes the new candidate and d is set to $\delta^-_{(N\setminus A)\cup\pi_{<i^*}}(i^*) = d^*$. If, on the other hand, both $\delta^-_{(N\setminus A)\cup\pi_{<i^*}\setminus\{i\}}(i^*) = d^* - 1$ and $\delta^-_{(N\setminus A)\cup\pi_{<i}}(i) = d^*$, then i stays the candidate and d remains equal to d^*. In both cases we have $d = d^*$, which implies that any future candidate j has at least d incoming edges from $(N \setminus A) \cup \pi_{<j}$. □

We now obtain our result for the 2-partition mechanism with permutation.

Theorem 4.3. *The 2-partition mechanism with permutation is impartial and* $7/12$*-optimal.*

Proof. Impartiality follows directly from Lemma 3.2 and Lemma 4.1.

Now consider a graph $G = (N, E)$, two distinct vertices $i_1^*, i_2^* \in N$ with $\delta^-(i_1^*) + \delta^-(i_2^*) = \Delta_2$, and let i_1 and i_2 be the two vertices selected by the mechanism from sets A_1 and A_2, respectively. We distinguish two cases, depending on whether i_1^* and i_2^* are in the same set or different sets of the partition (A_1, A_2).

First assume that i_1^* and i_2^* are in different sets, and without loss of generality that $i_1^* \in A_1$ and $i_2^* \in A_2$. In the permutation π used by the mechanism and chosen uniformly at random, an arbitrary vertex $i \in N \setminus \{i_1^*, i_2^*\}$ appears before or after each of i_1^* or i_2^* with equal probability, so

$$\mathbb{P}\left[i \in A_1 \cap \pi_{<i_1^*}\right] = \mathbb{P}\left[i \in A_1 \cap \bar{\pi}_{<i_1^*}\right] = \mathbb{P}\left[i \in A_2 \cap \pi_{<i_2^*}\right] = \mathbb{P}\left[i \in A_2 \cap \bar{\pi}_{<i_2^*}\right] = \frac{1}{4}.$$

When i_1^* is considered by the mechanism, so are any incoming edges from vertices in A_2 and any incoming edges from vertices in A_1 that appear in π before i_1^*. Thus, by Lemma 4.2,

$$\mathbb{E}\left[\delta^-(i_1)\right] \geq \mathbb{E}\left[\delta^-_{A_2\cup\pi_{<i_1^*}}(i_1^*)\right] = \sum\nolimits_{i\in N} \mathbb{P}\left[i \in A_2 \cup (A_1 \cap \pi_{<i_1^*})\right] \cdot \chi[(i, i_1^*) \in E],$$

where χ denotes the indicator function on Boolean expressions, i.e., $\chi[\phi] = 1$ if expression ϕ holds and $\chi[\phi] = 0$ otherwise. By taking i_2^* out of the sum and using that i_1^* and i_2^* are in different sets of

the partition and thus $\mathbb{P}[i_2^* \in A_2] = 1$, we obtain

$$\begin{aligned}
\mathbb{E}\left[\delta^-(i_1)\right] &\geq \sum\nolimits_{i \in N \setminus \{i_2^*\}} \Big(\mathbb{P}[i \in A_2 \cup (A_1 \cap \pi_{<i_1^*})] \cdot \chi[(i, i_1^*) \in E] \Big) \\
&\qquad + \mathbb{P}[i_2^* \in A_2 \cup (A_1 \cap \pi_{<i_1^*})] \cdot \chi[(i_2^*, i_1^*) \in E] \\
&= \sum\nolimits_{i \in N \setminus \{i_2^*\}} \Big((1 - \mathbb{P}[i \in (A_1 \cap \bar{\pi}_{<i_1^*})]) \cdot \chi[(i, i_1^*) \in E] \Big) \\
&\qquad + \chi[(i_2^*, i_1^*) \in E] \\
&= \sum\nolimits_{i \in N \setminus \{i_2^*\}} \Big(1 - \frac{1}{4} \Big) \chi[(i, i_1^*) \in E] + \chi[(i_2^*, i_1^*) \in E] \\
&\geq \frac{3}{4} \sum_{i \in N} \chi[(i, i_1^*) \in E] \\
&= \frac{3}{4} \delta^-(i_1^*).
\end{aligned}$$

As the same line of reasoning applies to i_2^*, we have $\mathbb{E}\left[\delta^-(i_2)\right] \geq \frac{3}{4}\delta^-(i_2^*)$ and conclude for this case that

$$\mathbb{E}\left[\frac{\delta^-(i_1, i_2)}{\Delta_2}\right] \geq \frac{\frac{3}{4}\delta^-(i_1^*) + \frac{3}{4}\delta^-(i_2^*)}{\delta^-(i_1^*) + \delta^-(i_2^*)} \geq \frac{3}{4}.$$

Now assume that i_1^* and i_2^* are in the same set of the partition, and without loss of generality that $i_1^*, i_2^* \in A_1$ and $\delta^-(i_1^*) \geq \delta^-(i_2^*)$. In the permutation π used by the mechanism and chosen uniformly at random, an arbitrary vertex $i \in N \setminus \{i_1^*, i_2^*\}$ appears before, between, or after i_1^* and i_2^* with probability 1/3 each, so

$$\begin{aligned}
\mathbb{P}\left[i \in A_2\right] &= \frac{1}{2} \quad \text{and} \\
\mathbb{P}\left[i \in A_1 \cap \pi_{<i_1^*} \cap \pi_{<i_2^*}\right] &= \mathbb{P}\left[i \in A_1 \cap ((\pi_{<i_1^*} \cap \bar{\pi}_{<i_2^*}) \cup (\bar{\pi}_{<i_1^*} \cap \pi_{<i_2^*}))\right] \\
&= \mathbb{P}\left[i \in A_1 \cap \bar{\pi}_{<i_1^*} \cap \bar{\pi}_{<i_2^*}\right] = \frac{1}{6}.
\end{aligned}$$

If $i_1^* \in \bar{\pi}_{<i_2^*}$, a possible edge from i_2^* to i_1^* would be considered by the mechanism, and by Lemma 4.2,

$$\begin{aligned}
\mathbb{E}\left[\delta^-(i_1)\right] &\geq \mathbb{E}\left[\delta^-_{A_2 \cup \pi_{<i_1^*}}(i_1^*)\right] \\
&\geq \sum\nolimits_{i \in N \setminus \{i_2^*\}} \Big(\mathbb{P}[i \in A_2 \cup (A_1 \cap \pi_{<i_1^*})] \cdot \chi[(i, i_1^*) \in E] \Big) \\
&\qquad + \mathbb{P}[i_2^* \in A_2 \cup (A_1 \cap \pi_{<i_1^*})] \cdot \chi[(i_2^*, i_1^*) \in E] \\
&= \sum\nolimits_{i \in N} \Big((1 - \mathbb{P}[i \in (A_1 \cap \bar{\pi}_{<i_1^*})]) \cdot \chi[(i, i_1^*) \in E] \Big) \\
&\qquad + \mathbb{P}[i_2^* \in (A_1 \cap \pi_{<i_1^*})] \cdot \chi[(i_2^*, i_1^*) \in E].
\end{aligned}$$

Since we assumed that $i_1^* \in \bar{\pi}_{<i_2^*}$, we have $\mathbb{P}\left[i \in (A_1 \cap \bar{\pi}_{<i_1^*})\right] = \mathbb{P}\left[i \in (A_1 \cap \bar{\pi}_{<i_1^*} \cap \bar{\pi}_{<i_2^*})\right] = 1/6$ and obtain

$$\mathbb{E}\left[\delta^-(i_1)\right] \geq \frac{5}{6} \sum_{i \in N} \chi[(i, i_1^*) \in E] = \frac{5}{6}\delta^-(i_1^*).$$

Analogously, if $i_2^* \in \bar{\pi}_{<i_1^*}$,

$$\mathbb{E}\left[\delta^-(i_1))\right] \geq \mathbb{E}\left[\delta^-_{A_2 \cup \pi_{<i_2^*}}(i_2^*)\right] = \frac{5}{6}\delta^-(i_2^*).$$

ALGORITHM 5: The randomized bidirectional permutation mechanism

Input: Graph $G = (N, E)$
Output: Set $\{i_1, i_2\} \subseteq N$ of at most two vertices
Choose a permutation $(\pi_1, \dots, \pi_n)$ of N uniformly at random;
Invoke Algorithm 1, the bidirectional permutation mechanism, for G and π

As each of the two events takes places with probability $1/2$, we conclude for this case that

$$\mathbb{E}\left[\frac{\delta^-(i_1 \cup i_2)}{\Delta_2}\right] \geq \frac{\frac{1}{2}\left(\frac{5}{6}\delta^-(i_1^*) + \frac{5}{6}\delta^-(i_2^*)\right)}{\delta^-(i_1^*) + \delta^-(i_2^*)} = \frac{5}{12}.$$

Averaging over both cases we finally obtain

$$\alpha \geq \frac{1}{2}\left(\frac{3}{4} + \frac{5}{12}\right) = \frac{7}{12},$$

as claimed. □

The 2-partition mechanism with permutation improves on the best deterministic mechanism for 2-selection, and it is natural to ask whether it can be improved upon further by a randomized 2-selection mechanism that is not exact. The answer to this question is not obvious: while the ability to select fewer vertices may make impartiality easier to achieve, actually selecting fewer vertices runs counter to the objective of selecting vertices with a large sum of indegrees. Indeed, in the case of 1-selection, no separation exists between exact and inexact mechanisms. For 2-selection, an obvious approach turns out to be effective: taking the best deterministic mechanism, which uses both directions of a fixed permutation, and invoking it for a random permutation. The resulting mechanism, which we call the randomized bidirectional permutation mechanism, is shown as Algorithm 5.

THEOREM 4.4. *The randomized bidirectional permutation mechanism is impartial and* $2/3$*-optimal.*

PROOF. The proof of Theorem 3.4 shows impartiality for any permutation that does not depend on the input to the mechanism, including one that is chosen uniformly at random.

Now consider a graph $G = (N, E)$, two distinct vertices $i_1^*, i_2^* \in N$ with $\delta^-(i_1^*) + \delta^-(i_2^*) = \Delta_2$, and let $i_1 = \Xi_\pi(G)$ and $i_2 = \Xi_{\bar{\pi}}(G)$ for the permutation π used by the mechanism. Assume without loss of generality that i_1^* appears before i_2^* in π, i.e., that $i_1^* \in \pi_{<i_2^*}$. As π was chosen uniformly at random, an arbitrary vertex $i \in N \setminus \{i_1^*, i_2^*\}$ appears before, between, or after i_1^* and i_2^* with probability $1/3$ each. By applying Lemma 3.3 to both i_1 and i_2,

$$\begin{aligned}\mathbb{E}\left[\frac{\delta^-(\{i_1, i_2\})}{\Delta_2}\right] &\geq \mathbb{E}\left[\frac{\delta^-_{\pi_{<i_1}}(i_1) + \delta^-_{\bar{\pi}_{<i_2}}(i_2)}{\Delta_2}\right] \\ &\geq \mathbb{E}\left[\frac{\max\left\{\delta^-_{\pi_{<i_1^*}}(i_1^*), \delta^-_{\pi_{<i_2^*}}(i_2^*)\right\} + \max\left\{\delta^-_{\bar{\pi}_{<i_1^*}}(i_1^*), \delta^-_{\bar{\pi}_{<i_2^*}}(i_2^*)\right\}}{\Delta_2}\right].\end{aligned} \tag{1}$$

Recall that possibly $i_1 = i_2$, and note that the bound is correct in this case as well. To bound the right-hand side of (1), we use the assumption that $i_1^* \in \pi_{<i_2^*}$ and observe that

$$\begin{aligned}
\mathbb{E}\left[\max\left\{\delta^-_{\pi_{<i_1^*}}(i_1^*), \delta^-_{\pi_{<i_2^*}}(i_2^*)\right\}\right] &\geq \mathbb{E}\left[\delta^-_{\pi_{<i_2^*}}(i_2^*)\right] \\
&= \sum_{i\in N} \mathbb{P}\left[i \in \pi_{<i_2^*}\right] \cdot \chi\left[(i, i_2^*) \in E\right] \\
&= \sum_{i\in N} (1 - \mathbb{P}\left[i \in \bar{\pi}_{<i_2^*}\right]) \cdot \chi\left[(i, i_2^*) \in E\right] \\
&= \frac{2}{3} \sum_{i\in N} \chi\left[(i, i_2^*) \in E\right] \\
&= \frac{2}{3}\, \delta^-(i_2^*).
\end{aligned}$$

Note that this bound only gets better when $(i_1^*, i_2^*) \in E$, as by assumption $\mathbb{P}[i_1^* \in \pi_{<i_2^*}] = 1$. An analogous argument for the other direction yields

$$\mathbb{E}\left[\max\left\{\delta^-_{\bar{\pi}_{<i_1^*}}(i_1^*), \delta^-_{\bar{\pi}_{<i_2^*}}(i_2^*)\right\}\right] \geq \frac{2}{3}\, \delta^-(i_1^*),$$

and by plugging both bounds into (1) we obtain

$$\mathbb{E}\left[\frac{\delta^-(i_1) + \delta^-(i_2)}{\Delta_2}\right] \geq \frac{\frac{2}{3}\delta^-(i_1^*) + \frac{2}{3}\delta^-(i_2^*)}{\delta^-(\{i_1^*, i_2^*\})} \geq \frac{2}{3},$$

as claimed. □

It is not hard to see that our analysis of the 2-partition mechanism with permutation and the bidirectional permutation mechanism is tight.

Theorem 4.5. *The 2-partition mechanism with permutation is at most 7/12-optimal. The randomized bidirectional permutation mechanism is at most 2/3-optimal.*

Proof. Consider a graph with a large number of vertices and only two edges (i, i_1^*) and (i, i_2^*), and observe that the maximum overall indegree of any set of two vertices is 2.

The 2-partition mechanism with permutation independently and uniformly at random assigns each of i_1^*, i_2^*, and i to one of two sets, such that in particular i_1^* and i_2^* are in the same set with probability 1/2 and in different sets with probability 1/2. If i_1^* and i_2^* are in the same set, the mechanism selects at most one of them. If i is in the respective other set, i.e., with probability 1/2, this happens with probability 1. If i is in the same set, it happens only with probability 2/3, namely when i appears before either i_1^* and i_2^* in a permutation π chosen uniformly at random. If i_1^* and i_2^* are in different sets, one of them is selected with probability 1, the other only when i appears before it in π, which happens with probability 1/2. In summary we thus expect $\alpha \leq \left(\frac{1}{2}(\frac{1}{2} + \frac{1}{2}\cdot\frac{2}{3}) + \frac{1}{2}(1+\frac{1}{2})\right)/2 = \frac{7}{12}$.

The randomized bidirectional permutation mechanism selects one of i_1^* and i_2^* with probability 1, the other only when i appears between i_1^* and i_2^* in a permutation π chosen uniformly at random, which happens with probability 1/3. Thus we conclude that $\alpha \leq \left(\frac{1}{3}\cdot 2 + \frac{2}{3}\cdot 1\right)/2 = \frac{2}{3}$. □

As special cases of Theorem 6.2 and Theorem 6.3 in Section 6, we will respectively obtain upper bounds of 3/4 and 2/3 for 2-selection mechanisms without and with exactness. These bounds suggest that neither the randomized bidirectional permutation mechanism nor the 2-partition mechanism with permutation is the best mechanism within its class. Figure 2 shows a 3/4-optimal impartial mechanism selecting at most two of three vertices, which certifies that the randomized

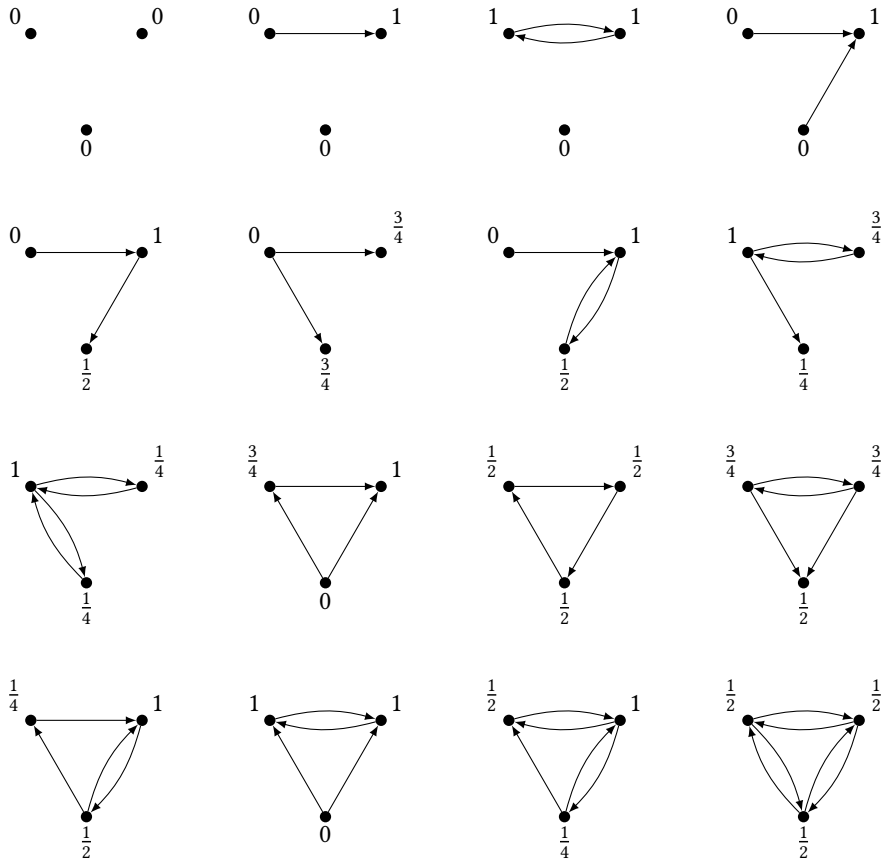

Fig. 2. A 3/4-optimal impartial mechanism for $n = 3$ and $k = 2$ given explicitly by the selection probabilities for all 16 voting graphs. The bound of 3/4 is best possible by Theorem 6.2.

bidirectional permutation mechanism is indeed not the best and that relaxing exactness is strictly beneficial.

The mechanism of Figure 2 can be obtained as the solution of an optimization problem to maximize the expected overall indegree of the vertices selected subject to impartiality, and its lack of universal impartiality illustrates one of the main obstacles that prevent us from obtaining tighter bounds and generalize our results to the selection of more than two vertices. Here, a mechanism is called universally impartial if it is a convex combination of deterministic impartial mechanisms. Mechanisms that are impartial but not universally impartial are notoriously difficult to analyze and sometimes exhibit rather peculiar behavior. The last two rows of the rightmost column of Figure 2 for example show a decrease in the probability of selecting two of the vertices as their indegrees go up, and this is both necessary for 3/4-optimality and difficult to justify.

5 SELECTING MORE THAN TWO AGENTS

The central component of our best inexact mechanisms, its use of one or both of the directions of a random permutation, does not generalize in any obvious way to the selection of additional vertices.

ALGORITHM 6: The k-partition mechanism with permutation

Input: Graph $G = (N, E)$ with $n \geq 2$
Output: Vertices $i_1, \ldots, i_k \in N$
Assign each $i \in N$ independently and uniformly at random to one of k sets $A_1, \ldots, A_k$;
Choose a permutation $(\pi_1, \ldots, \pi_n)$ of N uniformly at random;
for $j = 1, \ldots, k$ **do**
  $i_j := \Xi_{\pi, A_j}(G)$; ▷ `select one vertex from each set using extraction mechanism`
end
for $j = 1, \ldots, k$ **do**
  if $A_j = \emptyset$ **then**
    Choose i_j uniformly at random from $N \setminus \{i_1, \ldots, i_k\}$;
  end
end
return $\{i_1, \ldots, i_k\}$;

Our understanding of deterministic mechanisms for the selection of more than two vertices is particularly limited, but we can obtain a bound of $1/k$ by observing that the selection of only two instead of k vertices reduces the guarantee by a factor of at most $2/k$ and applying this observation to the bidirectional permutation mechanism.

A better bound for the randomized case, even with exactness, is achieved by a natural generalization of the 2-partition mechanism with permutation that uses a partition into k sets. The general mechanism is described formally as Algorithm 6. Its impartiality is easy to see, and we use an argument similar to that in the proof of Lemma 4.2 to obtain a performance guarantee that approaches $1 - 1/e$ as k grows.

THEOREM 5.1. *The k-partition mechanism with permutation is impartial and α-optimal for $\alpha = \frac{k}{k+1}\left(1 - (\frac{k-1}{k})^{k+1}\right)$.*

PROOF. Impartiality follows directly from Lemma 3.2 and Lemma 4.1.

Now consider a graph $G = (N, E)$ and a set I^* of vertices with $|I^*| = k$ and $\sum_{i \in I^*} \delta^-(i) = \Delta_k$, and denote the vertex selected by the mechanism from A_j by i_j, for $j = 1, \ldots, k$. For a fixed set A_j with $A_j \cap I^* \neq \emptyset$, let $i^* \in A_j \cap I^*$ such that $i \in \pi_{<i^*}$ for all $i \in A_j \cap I^* \setminus \{i^*\}$. Then,

$$\begin{aligned}\delta^-(i_j) &\geq \delta^-_{N \setminus A_j}(i^*) + \delta^-_{A_j \cap \pi_{<i^*}}(i^*) \\ &= \delta^-_{N \setminus (A_j \cup I^*)}(i^*) + \delta^-_{(A_j \cap \pi_{<i^*}) \setminus I^*}(i^*) + \delta^-_{I^*}(i^*),\end{aligned}$$

where the inequality holds by a similar argument as in the proof of Lemma 4.2 and the equality because we have chosen i^* to be the vertex in $A_j \cap I^*$ that appears last in π.

In the permutation π used by the mechanism and chosen uniformly at random, a given vertex appears after i^* with probability $|A_j \cap I^*|/(|A_j \cap I^*| + 1)$, so

$$\begin{aligned}&\mathbb{E}\left[\delta^-_{(A_j \cap \pi_{<i^*}) \setminus I^*}(i^*) \,\middle|\, |A_j \cap I^*| = l\right] \\ &\quad= \sum_{i \in N \setminus I^*} \mathbb{P}\left[i \in A_j \,\middle|\, |A_j \cap I^*| = l\right] \cdot \mathbb{P}\left[i \in \pi_{<i^*} \,\middle|\, |A_j \cap I^*| = l\right] \cdot \chi\left[(i, i^*) \in E\right] \\ &\quad= \frac{1}{k} \cdot \frac{l}{l+1}\left(\mathbb{E}\left[\delta^-(i^*) - \delta^-_{I^*}(i^*)\right]\right),\end{aligned}$$

where we have used that for $i \in N \setminus I^*$, $\mathbb{P}\left[i \in A_j \,\middle|\, |A_j \cap I^*| = l\right] = \mathbb{P}\left[i \in A_j\right] = 1/k$. Similarly,

$$\begin{aligned}\mathbb{E}\left[\delta^-_{N\setminus(A_j\cup I^*)}(i^*) \,\middle|\, |A_j \cap I^*| = l\right] &= \sum_{i\in N\setminus I^*} \mathbb{P}\left[i \in N \setminus A_j \,\middle|\, |A_j \cap I^*| = l\right] \cdot \chi\left[(i,i^*) \in E\right] \\ &= \frac{k-1}{k}\left(\mathbb{E}\left[\delta^-(i^*) - \delta^-_{I^*}(i^*)\right]\right).\end{aligned}$$

Thus

$$\begin{aligned}\mathbb{E}\left[\delta^-(i_j) \,\middle|\, |A_j \cap I^*| = l\right] &\geq \left(\frac{k-1}{k} + \frac{l}{k(l+1)}\right) \cdot \mathbb{E}\left[\delta^-(i^*) - \delta^-_{I^*}(i^*)\right] + \mathbb{E}\left[\delta^-_{I^*}(i^*)\right] \\ &\geq \left(\frac{k-1}{k} + \frac{l}{k(l+1)}\right)\frac{\Delta_k}{k},\end{aligned}$$

and by linearity of expectation,

$$\begin{aligned}\mathbb{E}\left[\frac{\sum_{j=1}^k \delta^-(i_j)}{\Delta_k}\right] &= \frac{1}{\Delta_k}\sum_{j=1}^k \mathbb{E}\left[\delta^-(i_j)\right] \\ &= \frac{1}{\Delta_k}\sum_{j=1}^k\sum_{l=1}^k \mathbb{E}\left[\delta^-(i_j) \,\middle|\, |I^* \cap A_j| = l\right] \cdot \mathbb{P}\left[|I^* \cap A_j| = l\right] \\ &\geq \frac{k}{\Delta_k}\sum_{l=1}^k \left(\frac{k-1}{k} + \frac{l}{k(l+1)}\right)\frac{\Delta_k}{k} \cdot \binom{k}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l} \\ &= \frac{k-1}{k}\sum_{l=1}^k \binom{k}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l} + \sum_{l=0}^k \frac{l}{k(l+1)}\binom{k}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l} \\ &= \frac{k-1}{k}\left(1-\left(1-\frac{1}{k}\right)^k\right) + \sum_{l=0}^k \frac{l}{k(l+1)}\binom{k}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l}. \qquad (2)\end{aligned}$$

We can now calculate

$$\begin{aligned}\sum_{l=0}^k \frac{l}{k(l+1)}\binom{k}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l} &= \sum_{l=0}^k \frac{1}{k}\left(1-\frac{1}{l+1}\right)\binom{k}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l} \\ &= \frac{1}{k} - \frac{1}{k}\sum_{l=0}^k \frac{1}{l+1}\cdot\frac{k!}{l!(k-l)!}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l} \\ &= \frac{1}{k} - \frac{1}{k}\sum_{l=0}^k \frac{1}{k+1}\binom{k+1}{l+1}k\left(\frac{1}{k}\right)^{l+1}\left(1-\frac{1}{k}\right)^{(k+1)-(l+1)} \\ &= \frac{1}{k} - \frac{1}{k+1}\sum_{l=1}^{k+1}\binom{k+1}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k+1-l} \\ &= \frac{1}{k} - \frac{1}{k+1}\left(1-\left(1-\frac{1}{k}\right)^{k+1}\right)\end{aligned}$$

and simplify (2) to conclude that

$$\mathbb{E}\left[\frac{\sum_{j=1}^{k}\delta^-(i_j)}{\Delta_k}\right] \geq \left(1-\frac{1}{k}\right)\left(1-\left(1-\frac{1}{k}\right)^k\right)+\frac{1}{k}-\frac{1}{k+1}\left(1-\left(1-\frac{1}{k}\right)^{k+1}\right)$$
$$= \frac{k}{k+1}\left(1-\left(\frac{k-1}{k}\right)^{k+1}\right),$$

as claimed. □

It is again not hard to see that this analysis is tight.

Theorem 5.2. *If the k-partition mechanism with permutation is α-optimal, then $\alpha \leq \frac{k}{k+1}\left(1-(\frac{k-1}{k})^{k+1}\right)$.*

Proof. Consider a graph with a large number of vertices and only k edges $(i, i_1^*), \ldots, (i, i_k^*)$. Let $I^* = \{i_1^*, \ldots, i_k^*\}$.

When partitioning the vertices into the sets $A_1,\ldots,A_k$, it is without loss of generality to assume that $i \in A_1$. For each $j \in \{2, \ldots, k\}$, a vertex with indegree 1 is selected from A_j if and only if $I^* \cap A_j \neq \emptyset$. This happens with probability $1-(\frac{k-1}{k})^k$, so by linearity of expectation the expected sum of indegrees of the vertices selected from $A_2 \cup \cdots \cup A_k$ is

$$(k-1)\left(1-\left(1-\frac{1}{k}\right)^k\right).$$

From A_1, a vertex with indegree 1 is selected if $A_1 \cap \bar{\pi}_{<i} \cap I^* \neq \emptyset$, and this condition is in fact necessary with probability going to 1 as the number of vertices with indegree 0 goes to infinity. For any $l \in \{0, \ldots, k\}$ we have $\mathbb{P}[|A_1 \cap I^*| = l] = \binom{k}{l}(1/k)^l(1-1/k)^{k-l}$ and $\mathbb{P}\left[A_1 \cap I^* \cap \bar{\pi}_{<i} \neq \emptyset \,\middle|\, |A_1 \cap I^*| = l\right] = l/(l+1)$. The probability of selecting a vertex with indegree 1 from $A_1 \cap I^*$ thus goes to

$$\sum_{l=0}^{k}\frac{l}{l+1}\binom{k}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l}$$

as the number of vertices goes to infinity.

The maximum overall indegree of any set of k vertices in the graph is k, so

$$\alpha \leq \frac{k-1}{k}\left(1-\left(1-\frac{1}{k}\right)^k\right)+\sum_{l=0}^{k}\frac{l}{k(l+1)}\binom{k}{l}\left(\frac{1}{k}\right)^l\left(1-\frac{1}{k}\right)^{k-l}.$$

This expression is equal to the lower bound in (2), and we conclude that the analysis in the proof of Theorem 5.1 is tight. □

6 UPPER BOUNDS

We conclude by giving upper bounds on the performance of impartial k-selection mechanisms for any value of k, and for both deterministic mechanisms and randomized mechanisms with and without exactness.

The first set of bounds applies to deterministic mechanisms and shows that the bidirectional permutation mechanism is the best deterministic mechanism for $k = 2$.

Theorem 6.1. *Consider a deterministic k-selection mechanism that is α-optimal on $\mathcal{G}_n$, where $k < n$. Then $\alpha \leq (k-1)/k$.*

Proof. Consider a graph $G = (V, E)$ with n vertices where $k + 1$ vertices are arranged in a directed cycle and the remaining vertices do not have any outgoing edges, i.e., $V = \{1, \dots, n\}$ and $E = \{(i, i + 1) : i = 1, \dots, k\} \cup \{(k + 1, 1)\}$. Denote by F the set of vertices selected from G by an arbitrary deterministic k-selection mechanism, and observe that there exists $i \in \{1, \dots, k + 1\} \setminus F$. Let $G' = (V, E \setminus (\{i\} \times V))$, and observe that by impartiality, the mechanism does not select i from G'. The mechanism thus selects at most $k - 1$ out of the k vertices with positive indegree in G' and cannot be more than $(k - 1)/k$-optimal. □

The next result result applies to randomized mechanisms without the requirement of exactness and shows that the mechanism of Figure 2 for the case when $k = 2$ and $n = 3$ is best possible.

Theorem 6.2. *Consider a k-selection mechanism that is impartial and α-optimal on $\mathcal{G}_n$, where $k < n$. Then*

$$\alpha \leq \begin{cases} \frac{1}{2} & \text{if } k = 1 \\ \frac{3}{4} & \text{if } k = 2 \\ \frac{2k}{2k+1} & \text{if } 3 \leq k = n - 1 \\ \frac{k+1}{k+2} & \text{otherwise.} \end{cases}$$

Proof. It is without loss of generality to assume that the k-selection mechanism is symmetric, i.e., that it assigns equal probabilities to indistinguishable vertices.

First assume that $k \leq n - 2$, and consider the two graphs on n vertices where $k + 2$ of the vertices have edges as in Figure 3 and the remaining vertices do not have any incoming or outgoing edges. It is easily verified that any symmetric impartial mechanism must assign probabilities as shown in Figure 3. In the graph on the left, the mechanism chooses a set of vertices with expected overall indegree $(k + 1)p_1$, while the maximum overall indegree for a set of k vertices is k, so

$$\alpha \leq \frac{(k + 1)p_1}{k}.$$

In the graph on the right, the mechanism chooses a set of vertices with expected overall indegree $2p_1 + 2kp_2$, while the maximum overall indegree for a set of k vertices is $2k$. Thus

$$\alpha \leq \frac{2p_1 + 2kp_2}{2k} = \frac{p_1}{k} + p_2 \leq \frac{p_1}{k} + \left(1 - \frac{2p_1}{k}\right) = 1 - \frac{p_1}{k}.$$

where the second inequality holds because $2p_1 + kp_2 \leq k$ by the graph on the right. In summary,

$$\alpha \leq \min\left\{\frac{(k + 1)p_1}{k}, 1 - \frac{p_1}{k}\right\} \leq \frac{k + 1}{k + 2},$$

where the second inequality holds because the minimum takes its maximum value when the two terms are equal.

Now assume that $k = n - 1$ and consider the two graphs on n vertices and edges as in Figure 3. Any symmetric impartial mechanism must again assign probabilities as shown. In the graph on the left, the mechanism chooses a set of vertices with expected overall indegree kp_1, while the maximum overall indegree for a set of k vertices is k, so

$$\alpha \leq p_1.$$

In the graph on the right, the mechanism chooses a set of vertices with expected overall indegree $2p_1 + 2(k - 1)p_2$, while the maximum overall indegree for a set of k vertices is $2(k - 1) + 1$. Thus

$$\alpha \leq \frac{2p_1 + 2(k - 1)p_2}{2(k - 1) + 1} \leq \frac{2p_1 + 2(k - 1)(\frac{k}{k-1} - \frac{2p_1}{k-1})}{2k - 1} = \frac{2k - 2p_1}{2k - 1},$$

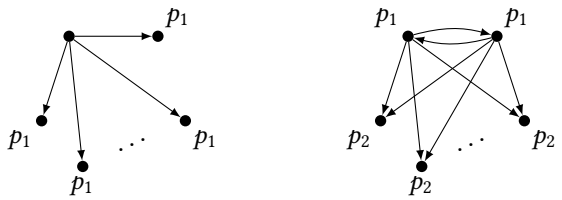

Fig. 3. Impartial probability assignment for two graphs with n vertices

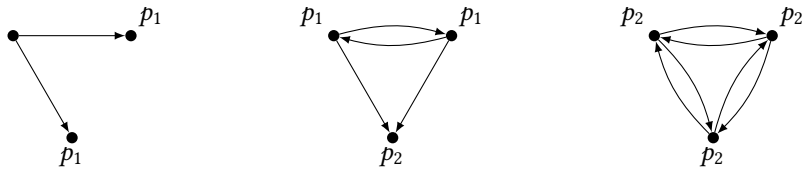

Fig. 4. Impartial probability assignment for three graphs with $n = 3$

where the second inequality holds because $2p_1 + (k-1)p_2 \leq k$ by the graph on the right. In summary,

$$\alpha \leq \min\left\{p_1, \frac{2k-2p_1}{2k-1}\right\} \leq \frac{2k}{2k+1},$$

where the second inequality holds because the minimum takes its maximum value when the two terms are equal.

In the special case where $k = 2$, an additional graph can be used to obtain a stronger bound. For this, consider situations where 3 vertices have outgoing edges as in Figure 4 and the remaining $n-3$ vertices do not have any outgoing edges. Note that the first two graphs are the same as those in Figure 3 when $k = 2$. It is again easily verified that any impartial mechanism must assign probabilities as shown. Thus

$$\alpha \leq \min\left\{\frac{2p_1}{2}, \frac{6p_2}{4}\right\} \leq \{p_1, 3-3p_1\} \leq \frac{3}{4},$$

where the first inequality holds by the first and third graph, the second inequality because $2p_1 + p_2 \leq 2$ by the second graph, and the third inequality because the minimum takes its maximum value when the two terms are equal.

The bound for $k = 1$ is easily obtained by considering the special case of the graphs in Figure 3 where two vertices have outgoing edges as shown and the others do not have any outgoing edges. Then

$$\alpha \leq p_1 \leq 1/2,$$

where the inequalities hold respectively by the first and second graph. □

Our final result concerns randomized mechanisms that are exact. It certifies that the 2-partition mechanism with permutation is best possible within this class when $k = 2$ and $n = 3$, and together with the mechanism of Figure 2 shows a strict separation between randomized mechanisms with and without exactness. It does not preclude improvements over the 2-partition mechanism with permutation when $n > 3$. A comparison with Theorem 6.2 further suggests that the influence of the exactness constraint may be limited to cases where almost all vertices are selected.

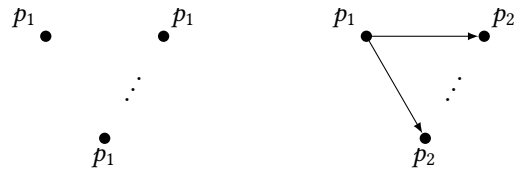

Fig. 5. Impartial probability assignment for two graphs with n vertices

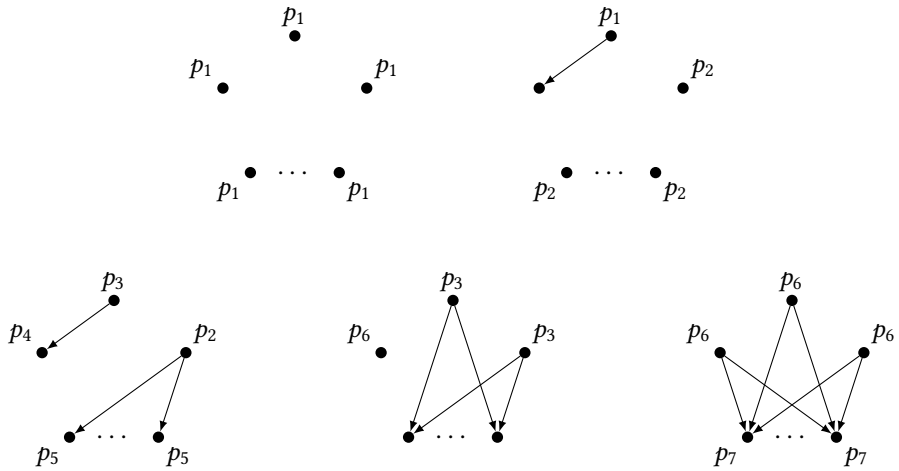

Fig. 6. Impartial probability assignment for five graphs with n vertices

Theorem 6.3. *Consider a k-selection mechanism that is exact, impartial, and α-optimal on $\mathcal{G}_n$, where $k < n$. Then*

$$\alpha \leq \begin{cases} \frac{1}{2} & \text{if } k = 1 \\ \frac{k}{k+1} & \text{if } 2 \leq k = n-1 \\ \frac{5}{7} & \text{if } 2 = k = n-2 \\ \frac{7k^3+5k^2-6k+12}{7k^3+13k^2-2k} & \text{if } 3 \leq k = n-2 \\ \frac{k+1}{k+2} & \text{otherwise.} \end{cases}$$

Proof. First assume that $2 \leq k = n-1$, and consider the two graphs with $n = k+1$ vertices shown in Figure 5. By impartiality, the probability of selecting the vertex at the top left must be equal for both graphs. Any symmetric mechanism assigns equal probabilities to all vertices in the left graph and equal probabilities to all vertices with indegree 1 in the right graph. Denoting the former probability by p_1 and the latter by p_2, exactness implies that $(k+1)p_1 = p_1 + kp_2 = k$, so $p_1 = k/(k+1)$, $p_2 = (k - k/(k+1))/k = k/(k+1)$, and thus $\alpha \leq p_2 = k/(k+1)$.

Now assume that $3 \leq k = n-2$, and consider the five graphs with five vertices shown in Figure 6. Using similar arguments as above it is easily established that any symmetric impartial mechanism must assign probabilities as shown. By the first and second graph, $p_1 = k/(k+2)$ and

$p_1 + kp_2 \geq k-1$, and thus

$$p_2 \geq \frac{k-1}{k} - \frac{p_1}{k} = \frac{k-1}{k} - \frac{1}{k+2} = \frac{(k-1)(k+2)-k}{k(k+2)}. \tag{3}$$

By the third graph, $p_2 + p_3 + p_4 + (k-1)p_5 = k$ and thus

$$\frac{1}{k}p_4 + \frac{k-1}{k}p_5 = 1 - \frac{1}{k}p_2 - \frac{1}{k}p_3 \leq 1 - \frac{(k-1)(k+2)-k}{k^2(k+2)} - \frac{1}{k}p_3 = \frac{k^3+k^2+2}{k^2(k+2)} - \frac{1}{k}p_3, \tag{4}$$

where the inequality holds by (3). By the fourth and fifth graph, $2p_3 + p_6 \geq 1$ and $3p_6 + (k-1)p_7 = k$, and thus

$$p_7 = \frac{k}{k-1} - \frac{3}{k-1}p_6 \leq \frac{k}{k-1} - \frac{3}{k-1}(1-2p_3) = \frac{1}{k-1}(6p_3 + k - 3). \tag{5}$$

Finally, by the third and fifth graph,

$$\alpha \leq \min\left\{\frac{1}{k}p_4 + \frac{k-1}{k}p_5, p_7\right\} \leq \min\left\{\frac{k^3+k^2+2}{k^2(k+2)} - \frac{1}{k}p_3, \frac{1}{k-1}(6p_3+k-3)\right\}, \tag{6}$$

where the second inequality holds by (4) and (5). The minimum takes its maximum value when the two terms are equal, i.e., when

$$\begin{aligned}
\frac{k^3+k^2+2}{k^2(k+2)} - \frac{1}{k}p_3 &= \frac{1}{k-1}(6p_3+k-3),\\
\frac{6}{k-1}p_3 + \frac{1}{k}p_3 &= \frac{k^3+k^2+2}{k^2(k+2)} - \frac{k-3}{k-1},\\
\frac{7k-1}{k(k-1)}p_3 &= \frac{(k^3+k^2+2)(k-1) - k^2(k+2)(k-3)}{k^2(k+2)(k-1)}.
\end{aligned}$$

Thus

$$\begin{aligned}
p_3 &= \frac{(k^3+k^2+2)(k-1) - k^2(k+2)(k-3)}{k(k+2)(7k-1)}\\
&= \frac{k^4+k^3+2k-k^3-k^2-2-k^4+k^3+6k^2}{k(k+2)(7k-1)}\\
&= \frac{k^3+5k^2+2k-2}{k(k+2)(7k-1)},
\end{aligned}$$

and by plugging this into (6),

$$\begin{aligned}
\alpha &\leq \frac{k^3+k^2+2}{k^2(k+2)} - \frac{k^3+5k^2+2k-2}{k^2(k+2)(7k-1)}\\
&= \frac{7k^4+7k^3+14k-k^3-k^2-2-k^3-5k^2-2k+2}{k^2(k+2)(7k-1)}\\
&= \frac{7k^3+5k^2-6k+12}{k(k+2)(7k-1)}.
\end{aligned}$$

In the special case where $2 = k = n-2$, only a single vertex is chosen with probability p_5 in the third graph of Figure 6, so by symmetry $p_2 = p_3$ and $p_4 = p_5$. Thus

$$p_7 = 2 - 3p_6 \leq 2 - 3(1-2p_3) = 6p_3 - 1 = (6 - 6p_4) - 1 = 5 - 6p_4,$$

and

$$\alpha \leq \min\{p_4, p_7\} \leq \min\{p_4, 5-6p_4\} \leq \frac{5}{7},$$

where the last inequality holds because the minimum again takes its maximum value when the two terms are equal.

The bounds for the remaining two cases, of $1/2$ if $k = 1$ and of $(k + 1)/(k + 2)$ otherwise, follow directly from Theorem 6.2. □

REFERENCES

[1] N. Alon, F. Fischer, A. D. Procaccia, and M. Tennenholtz. 2011. Sum of Us: Strategyproof Selection from the Selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*. 101–110.

[2] H. Aziz, O. Lev, N. Mattei, J.S. Rosenschein, and T. Walsh. 2016. Strategyproof Peer Selection: Mechanisms, Analyses, and Experiments. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*. 390–396.

[3] G. Birkhoff. 1946. Tres Observaciones Sobre el Algebra Lineal. *Revista Facultad de Ciencias Exactas, Puras y Aplicadas Universidad Nacional de Tucumán, Serie A* 5 (1946), 147–151.

[4] N. Bousquet, S. Norin, and A. Vetta. 2014. A Near-Optimal Mechanism for Impartial Selection. In *Proceedings of the 10th International Conference on Web and Internet Economics*. 133–146.

[5] I. Caragiannis, A. Filos-Ratsikas, S.K.S. Frederiksen, K.A. Hansen, and Z. Tan. 2016. Truthful Facility Assignment with Resource Augmentation: An Exact Analysis of Serial Dictatorship. In *Proceedings of the 12th International Conference on Web and Internet Economics*. 236–259.

[6] G. de Clippel, H. Moulin, and N. Tideman. 2008. Impartial Division of a Dollar. *Journal of Economic Theory* 139, 1 (2008), 176–191.

[7] F. Fischer and M. Klimm. 2015. Optimal Impartial Selection. *SIAM J. Comput.* 44, 5 (2015), 1263–1285.

[8] R. Holzman and H. Moulin. 2013. Impartial Nominations for a Prize. *Econometrica* 81, 1 (2013), 173–196.

[9] D. Kurokawa, O. Lev, J. Morgenstern, and A.D. Procaccia. 2015. Impartial Peer Review. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*. 582–588.

[10] A. Mackenzie. 2015. Symmetry and Impartial Lotteries. *Games and Economic Behavior* 94 (2015), 15–28.

[11] A. Mackenzie. 2016. A Game of the Throne of Saint Peter. (2016). Working paper.

[12] M. Mitzenmacher, A.W. Richa, and R. Sitaraman. 2001. The Power of Two Random Choices: A Survey of Techniques and Results. In *Handbook of Randomized Computing*, S. Rajasekaran, P.M. Pardalos, J.H. Reif, and J. Rolim (Eds.). Vol. 1. Springer, 255–312.

[13] A. D. Procaccia and M. Tennenholtz. 2013. Approximate Mechanism Design Without Money. *ACM Transactions on Economics and Computation* 1, 4 (2013), Article 18.

[14] D.D. Sleator and R.E. Tarjan. 1985. Amortized Efficiency of List Update and Paging Rules. *Commun. ACM* 28, 2 (1985), 202–208.

[15] S. Tamura. 2016. Characterizing Minimal Impartial Rules for Awarding Prizes. *Games and Economic Behavior* 95 (2016), 41–46.

[16] S. Tamura and S. Ohseto. 2014. Impartial Nomination Correspondences. *Social Choice and Welfare* 43, 1 (2014), 47–54.

[17] J. Witkowski, Y. Bachrach, P. Key, and D. C. Parkes. 2013. Dwelling on the Negative: Incentivizing Effort in Peer Prediction. In *Proceedings of the 1st AAAI Conference on Human Computation and Crowdsourcing*. 190–197.

[18] J. Witkowski and D. C. Parkes. 2012. Peer Prediction without a Common Prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 964–981.