Strategy-Proofness of Scoring Allocation Correspondences

Nhan-Tam Nguyen, Dorothea Baumeister, and Jörg Rothe

Abstract

We study resource allocation in a model due to Brams and King [2005] and further developed by Baumeister et al. [2014]. Resource allocation deals with the distribution of resources to agents. We assume resources to be indivisible, nonshareable, and of single-unit type. Agents have ordinal preferences over single resources. Using scoring vectors, every ordinal preference induces a utility function. These utility functions are used in conjunction with utilitarian social welfare to assess the quality of allocations of resources to agents. Then allocation correspondences determine the optimal allocations that maximize utilitarian social welfare.

Since agents may have an incentive to misreport their true preferences, the question of strategy-proofness is important to resource allocation. We assume that a manipulator has responsive preferences over the power set of the resources. We use extension principles (from social choice theory, such as the Kelly and the Gärdenfors extension) for preferences to study manipulation of allocation correspondences. We characterize strategy-proofness of the utilitarian allocation correspondence: It is Gärdenfors/Kelly-strategy-proof if and only if the number of different values in the scoring vector is at most two or the number of occurrences of the greatest value in the scoring vector is larger than half the number of goods.

1 Introduction

Resource allocation deals with the distribution of scarce *resources* among *agents*, who may have different *preferences* over (subsets, also called *bundles*, of) resources. The goal is to find *allocations*, assignments of resources to agents, which satisfy certain criteria.

We use the model proposed by Brams and King [2005] (and followed up by Baumeister et al. [2014]) that reconciles the need of collective utility functions for interpersonal comparability with ordinal preferences, using scoring vectors to map ordinal preferences to additive utility functions. This mapping is performed as if there were the same differences in preferential intensity between resources across all agents. We assume that resources are indivisible, nonshareable, and of single-unit type, that is, they can neither be split nor belong to multiple agents simultaneously. This ordinal-preference model is motivated by the issue with cardinal preferences concerning elicitation and (to a certain extent) incomparability between two utility functions. It is far easier to ask for a ranking of (bundles of) resources than for a utility function where every resource/bundle has a numerical value. The tradeoff between ordinal and cardinal preferences has been studied by Caragiannis and Procaccia [2011] in voting, who show that the induced social welfare distortion is low.

Our main contribution is a characterization of strategy-proofness of the utilitarian allocation correspondence: No agent can benefit from misreporting her true preferences if and only if the number of different values in the scoring vector is at most two or the number of occurrences of the greatest value in the scoring vector is larger than the sum of the other values' numbers of occurrences (equivalently, the second part can be restated as "the number of occurrences of the greatest value is larger than half the number of goods").

In Section 2, the model for resource allocation with ordinal preferences due to Brams and King [2005] is presented (in the notation of Baumeister et al. [2014]). Then, in Section 3, strategy-proofness under that model is studied and the main result is shown: a characterization of strategy-proofness of the utilitarian allocation correspondence. Related work is described in Section 4. At last, Section 5 gives a summary and concludes with open problems and directions of future research.

2 Scoring Allocation Correspondences

Let $A = \{1, \dots, n\}$ be a set of *agents* and let $R = \{r_1, \dots, r_m\}$ be a set of indivisible, nonshareable *resources* (or *goods* or *objects* or *items*) of single-unit type. An *allocation* of resources to agents is given by a partition $\pi = (\pi_1, \dots, \pi_n)$, where $\pi_i \subseteq R$ is the bundle of resources assigned to agent i. Agents are assumed to have (ordinal) preferences over all bundles of resources. However, to ensure feasibility in practice, instead of ranking all subsets of R, we assume agents to rank only single resources. This is a crucial assumption: While it avoids a heavy elicitation burden and allows specifying our problems compactly, agents will not be able to express preferential dependencies between resources. Under this assumption, we consider a *preference profile* $P = (>_1, \dots, >_n)$ as a collection of n linear rankings over R. We define a *scoring allocation correspondence* as a function that maps any such profile to a nonempty subset of allocations.

Definition 1 *A* scoring vector $s = (s_1, \dots, s_m)$ *consists of real numbers satisfying* $s_1 \geq \cdots \geq s_m \geq 0$ *and* $s_1 > 0$*. For a preference ranking* $>$ *over* R *and a resource* $r \in R$*, denote the* rank of r under $>$ *by* $rank(r, >)$*. For each bundle* $B \subseteq R$*, define the* utility function over 2^R induced by the ranking $>$ on R and the scoring vector s *by* $u_{>,s}(B) = \sum_{r \in B} s_{rank(r,>)}$.

Bouveret and Lang [2011] consider the following specific scoring vectors for allocating indivisible goods:

- Borda scoring: $\text{borda} = (m, m-1, \dots, 1)$,[1]
- lexicographic scoring: $\text{lex} = (2^{m-1}, 2^{m-2}, \dots, 1)$,
- quasi-indifference for some ε, $0 < \varepsilon \ll 1$:
  $\varepsilon\text{-qi} = (1 + (m-1)\varepsilon, 1 + (m-2)\varepsilon, \dots, 1)$.

In addition, we also consider the following scoring vector:

- k-approval: $k\text{-app} = (1, \dots, 1, 0, \dots, 0)$, where the first k entries are all ones and the remaining entries are zeros.

A monotonic, symmetric *aggregation function* can then be used to aggregate the individual utilities with the goal of maximizing the overall utility. Typically, one considers utilitarianism (i.e., the *sum*), and two versions of egalitarianism (namely, *min* and *leximin*); see, e.g., Baumeister et al. [2014]. We restrict ourselves to utilitarian social welfare maximization: $F_s(P) = \text{argmax}_\pi \sum_{1 \leq i \leq n} u_{>_i,s}(\pi_i)$, for preference profile $P = (>_1, \dots, >_n)$ and allocation $\pi = (\pi_1, \dots, \pi_n)$. From now on we slightly abuse notation by writing $u_{>_i,s}(\pi)$ instead of $u_{>_i,s}(\pi_i)$ to denote the utility for agent i under allocation $\pi = (\pi_1, \dots, \pi_n)$.

3 Strategy-Proofness

Although the central authority, which collects the rankings and computes the output, knows only the agents' rankings over singletons, the manipulating agent has full knowledge of her own preferences over all sets of objects. We make the crucial assumptions that every good has a positive value and that the agents' (unrevealed) strict preference relations are responsive,[2] which justifies the fact that they each can "safely" project their preference relation over 2^R onto a preference relation over R.

[1] The common definition of Borda scoring in voting is based on the vector $(m-1, m-2, \dots, 1, 0)$. However, we follow Brams et al. [2003] by setting the score of the bottom-rank resource to the value 1. Note that scoring vectors in voting can be shifted or scaled without changing the winner set [Hemaspaandra and Hemaspaandra, 2007]; but for scoring allocation correspondences such an operation would have an impact in general [Baumeister et al., 2014].

[2] $\succ$ is *responsive* if for all $X \subseteq R$, $x \in X$, and $x' \in R \setminus X$, we have $X \succ X \setminus \{x\}$ and if $\{x\} \succ \{x'\}$ then $X \succ (X \setminus \{x\}) \cup \{x'\}$.

This is also consistent with the common assumption of free disposal and that scoring vectors are positive.

Turning now to strategy-proofness for scoring allocation *correspondences* F_s we have to compare *sets* of winning allocations that can emerge from an original profile P and a modified profile P'. We assume that the final winning allocation is chosen by an unknown random device which assigns positive probability to each winning allocation. For this we make use of *extension principles*, which are common in social choice theory. An extension principle lifts a preference relation over single alternatives to a preference relation over sets of alternatives. Here, from the point of view of agent i, alternatives are possible shares, i.e., sets of objects. We consider the two classical extension principles that are due to Kelly [1977] and Gärdenfors [1976].[3] Note that the extension principles that we use require that agents are pessimistic: No manipulation may result in a worse outcome, which means that agents assume that the worst possible allocation is chosen. Furthermore, these extension principles assume that alternatives are mutually exclusive: Only one of the alternatives among the winners is ultimately chosen. This choice is performed under *complete uncertainty*. However, bundles (alternatives) are not mutually exclusive in our model of manipulation. Suppose that we have two bundles A, B with $A \subset B$. Then A and B are *not* mutually exclusive. Nevertheless, the pessimistic manipulator can pretend that the bundles are mutually exclusive for the following reason: Even if the manipulator knows that a specific subset of a bundle is realized in every allocation, a pessimistic manipulator will misreport her preferences only if this does not result in a worse outcome for her. Since the probabilities that govern the random device are not known, two sets of winning bundles, $F(P)$ and $F(P')$, can have completely different probabilities attached to each bundle. For example, if $F(P')$ extends every bundle in $F(P)$ by an additional resource and thus, in addition, may intuitively seem to be more appealing than $F(P)$ because an additional resource is guaranteed, the best bundle (according to the manipulator's preferences) could be assigned a very low probability in $F(P')$, whereas in $F(P)$ it could have a very high probability. Because we assume that manipulators are pessimistic, this is also an argument against the use of, e.g., an injective mapping that assigns to every alternative a in $F(P)$ an alternative b in $F(P')$ such that b is preferred to a.

To define extensions, it is easier to start with a weak order $\geq$ associated with a strict linear order $>$, where $x \geq y$ if and only if $x > y$ or $x = y$ (here, "$=$" expresses indifference).

Definition 2 *Let $\geq$ be a weak order over R and $A, B \subseteq R$. Define the* Kelly extension of $\geq$ *by $A \succeq^K B$ if and only if for all $x \in A$ and $y \in B$ we have $x \geq y$. Define the* Gärdenfors extension of $\geq$ *by $A \succeq^G B$ if and only if one of the following conditions holds:*

(a) $A \subset B$, and for all $x \in A$ and $y \in B \setminus A$, we have $x \geq y$;

(b) $B \subset A$, and for all $x \in A \setminus B$ and $y \in B$, we have $x \geq y$;

(c) neither $A \subset B$ nor $B \subset A$, and for all $x \in A \setminus B$ and $y \in B \setminus A$, we have $x \geq y$.

Finally, $A \succ^K B$ if and only if $A \succeq^K B$ and not $B \succeq^K A$, and similarly for $\succ^G$.

The definitions of $\succeq^K$ and $\succeq^G$ naturally carry over from sets of resources to sets of allocations of resources.

We now define manipulation of allocation correspondences according to these extension principles as follows.

Definition 3 *Let $e \in \{K, G\}$ be an extension principle, F an allocation correspondence, $P = (>_1, \dots, >_n)$ a preference profile, $P' = (>'_i, >_{-i})$ the profile identical to P except that agent i's preference $>_i$ is changed to $>'_i$, and let $\succ_i$ be a responsive preference relation over 2^R extending $>_i$. Define $F(P)_i = \{\pi_i \mid \pi \in F(P)\}$.*

[3] We don't explicitly consider the extension principle proposed by Fishburn [1972], which is intermediate between Kelly's and Gärdenfors's; note that our results therefore apply to it as well.

We say that $>'_i$ *is an e*-manipulation for P, F, and $\succ_i$ *if* $F(P')_i \succ^e_i F(P)_i$. *A scoring allocation correspondence* F *is e*-manipulable by an agent i *if there exist* $P = (>_1, \dots, >_n)$, $\succ_i$ *extending* $>_i$, *and* $>'_i$ *such that* $>'_i$ *is an e-manipulation for profile* P, F, *and* $\succ_i$. F *is e*-strategy-proof *if it is not e-manipulable by any agent.*

One can show that for each agent i, $F(P')_i \succ^K_i F(P)_i$ implies $F(P')_i \succ^G_i F(P)_i$. Therefore, Kelly-manipulability implies Gärdenfors-manipulability; equivalently, Gärdenfors-strategy-proofness implies Kelly-strategy-proofness. These notions have been investigated so far only for irresolute social choice functions in the context of voting (see, e.g., the papers by Brandt [2015] and Brandt and Brill [2011]). We will apply them to scoring allocation correspondences.

Note that one cannot define manipulation in terms of the agents' scores (induced utility functions) because their scores do not necessarily reflect the agents' preferences. Recall that scoring allocation correspondences merely use scores as proxies to "approximate" agents' preferences.

To summarize our model, agents have preferences over the power set of objects. However, they only submit a linear order over the objects to the scoring allocation procedure. The scoring allocation procedure uses a scoring vector to cardinalize the ordinal preferences and then optimizes social welfare with respect to the proxy utility functions. Since optimal allocations are not necessarily unique, agents lift their preferences over bundles of objects to preferences over sets of bundles of objects. These lifted preferences are used to determine whether a manipulation is successful.

The following notation will be used heavily in the proofs of this paper.

Definition 4 $R_i(P) = \{r \in R \mid (\exists \pi_i \in F_s(P)_i)\,[r \in \pi_i]\}$ *is the set of all resources that agent i gets in some allocation of* $F_s(P)$.

A resource $r \in R_i(P)$ *is* sure in profile P for agent i *if* $r \in \pi_i$ *for all* $\pi_i \in F_s(P)_i$; $r \in R_i(P)$ *is* contested in P for i *if it is not sure in P for i. Let* $S_i(P)$ *denote the set of all sure resources in P for i.*

Definition 5 *Let* $s = (s_1, \dots, s_m)$ *be a scoring vector with k different values* $v_1, \dots, v_k$ *with* $v_1 > v_2 > \dots > v_k$. *Denote by* α_i *the number of occurrences of the i-th value among* $v_1, \dots, v_k$ *in s, that is,* $\alpha_i = \|\{j \mid v_i = s_j, 1 \le j \le m\}\|$. *The set of resources that are in the i-th bin for agent j with linear order* $>_j$ *is defined by* $A_i(>_j) = \{r \in R \mid 1 + \sum_{k=1}^{i-1} \alpha_k \le rank(r, >_j) \le \sum_{k=1}^{i} \alpha_k\}$. *We omit the scoring vector in all definitions when it is clear from context.*

Notation 6 *Let* $>$ *be a linear order on the set of resources* $R = \{r_1, \dots, r_m\}$, *say* $r_1 > r_2 > \dots > r_m$. *Let s be a scoring vector with k different values. We sometimes write* $>$ *as*

$$r_1\, r_2 \cdots r_{\alpha_1 - 1}\, r_{\alpha_1} \mid r_{\alpha_1+1} \cdots r_{\alpha_1+\alpha_2} \mid \cdots \mid r_{1+\sum_{i=1}^{k-1}\alpha_i} \cdots r_{\sum_{i=1}^{k}\alpha_i}.$$

Now we show Gärdenfors-strategy-proofness for the case of at most two different values in the scoring vector. The intuition behind the following proof is that swapping the positions of a high-ranked and a low-ranked resource always leads to a situation where the high-ranked resource is replaced by the low-ranked resource in a bundle of an optimal allocation.

Proposition 7 *For at least two agents and each scoring vector s having at most two different values,* F_s *is Gärdenfors-strategy-proof (and thus also Kelly-strategy-proof).*

Proof. Let R be the set of resources with $a, b \in R$, and let $P = (>_1, \dots, >_n)$, $n \ge 2$, be the profile where in the ranking of the last agent (the manipulator) a gets a higher score than b, that is, $a >_n b$ and we have two distinct score values, $\alpha = s_{rank(a,>_n)}$ and $\beta = s_{rank(b,>_n)}$ with $\alpha > \beta$. We distinguish the following four cases regarding the positions of resources a and b in the preferences of the remaining agents:

Case 1(a) There is some $i < n$ with $s_{rank(a,>_i)} = \alpha$ and there is some $j < n$ with $s_{rank(b,>_j)} = \alpha$. Then a is contested for agent n, but agent n never receives b.

Case 1(b): There is some $i < n$ with $s_{rank(a,>_i)} = \alpha$ and for all j, $1 \leq j < n$, $s_{rank(b,>_j)} = \beta$. Then a and b are contested for n.

Case 2(a): For all i, $1 \leq i < n$, we have $s_{rank(a,>_i)} = \beta$ and there is some $j < n$ with $s_{rank(b,>_j)} = \alpha$. Then a is sure for n, but agent n never receives b.

Case 2(b): For all i, $1 \leq i < n$, we have $s_{rank(a,>_i)} = \beta$ and for all j, $1 \leq j < n$, $s_{rank(b,>_j)} = \beta$. Then a is sure and b is contested for n.

Now, let $P'_n = (>_1, \ldots, >_{n-1}, >'_n)$ be the modified profile where the only change is that in the ranking of the last agent the positions of resources a and b are swapped. Going through the cases with the modified profile, we see that either (1) agent n never receives a in $F_s(P')$ or (2) agent n does not receive a in some $\pi \in F_s(P')$ and either (a) b is contested for n in P' or (b) b is sure for n in P'. Hence we have the above four cases 1(a), 1(b), 2(a), and 2(b). In all of them it holds that for each $\pi \in F_s(P)$ with $a \in \pi_n$ and $b \notin \pi_n$ (which exists in all cases), there is some $\pi' \in F_s(P')$ with $\pi'_n = (\pi_n \setminus \{a\}) \cup \{b\}$. By responsiveness of $\succ_n$ and since $a >_n b$, $\pi_n \succ_n \pi'_n$ follows. Now we show that we can use the sets π_n and π'_n in all cases of the Gärdenfors extension:

Case I: $F_s(P')_n \subset F_s(P)_n$. This is Case 1(b). It holds that π'_n is in $F_s(P')_n$ and π_n is in $F_s(P)_n \setminus F_s(P')_n$ because of $b \notin \pi_n$ but $b \in \pi''_n$ is true for all $\pi''_n \in F_s(P')_n$.

Case II: $F_s(P)_n \subset F_s(P')_n$. This is Case 2(a). It holds that π_n is in $F_s(P)_n$ and π'_n is in $F_s(P')_n \setminus F_s(P)_n$ because of $b \in \pi'_n$ but $b \notin \pi''_n$ is true for all $\pi''_n \in F_s(P)_n$.

Case III: Neither $F_s(P')_n \subset F_s(P)_n$ nor $F_s(P)_n \subset F_s(P')_n$. These are Cases 1(a) and 2(b). For both of them, π_n is in $F_s(P)_n \setminus F_s(P')_n$ and π'_n is in $F_s(P')_n \setminus F_s(P)_n$ with the same arguments as above.

Thus it is not true that $F_s(P'_n)_n \succ^G_n F_s(P)_n$. Since this argument analogously works for all agents, F_s is not G-manipulable by any agent, which means it is Gärdenfors-strategy-proof. ❑

The idea of the next result is that if there are at least three different values and the first bin is not large enough, then all resources in the first bin of some agent fit into lower ranked bins of some other agent. This can be used to "shift" resources into the right bins.

Proposition 8 *Let s be a scoring vector with $k \geq 3$ different values satisfying $\alpha_1 \leq \sum_{i=2}^{k} \alpha_i$. Then F_s is K-manipulable.*

Proof. The goal is to construct profiles showing that F_s is K-manipulable by some agent. Let $R = \{r_1, \ldots, r_m\}$ be the set of resources. We distinguish the following two cases.

Case 1: $\alpha_1 = \sum_{i=2}^{k} \alpha_i$. Let profile P consist of two linear orders, where agent 2 is the manipulator:

$$r_1 >_1 r_2 >_1 \cdots >_1 r_{\alpha_1} >_1 \cdots >_1 r_m$$
$$r_{\alpha_1+1} >_2 r_{\alpha_1+2} >_2 \cdots >_2 r_m >_2 r_1 >_2 r_2 >_2 \cdots >_2 r_{\alpha_1}$$

We have $S_2(P) = \{r_{\alpha_1+1}, r_{\alpha_1+2}, \ldots, r_m\}$. Because none of the resources in $A_1(>_1)$ are assigned to agent 2, $S_2(P)$ is the only bundle that agent 2 receives in an allocation. The manipulator's goal is to receive an additional resource without losing one of her sure resources. This can be achieved by swapping resource $r_{\alpha_1+\alpha_2+1}$ with resource r_1 in $>_2$. Resource r_1 becomes contested because r_1 is in the first bin for both agents. But resource $r_{\alpha_1+\alpha_2+1}$ is still sure for agent 2 in profile P' because there are at least three different values (that is, bins), and resource $r_{\alpha_1+\alpha_2+1}$ is in the third bin of agent 1 but it is in the second bin of agent 2's manipulation. Hence, agent 2 receives either $S_2(P)$ or $S_2(P) \cup \{r_1\}$ in an allocation. We thus have a K-manipulation by agent 2 as required.

Case 2: $\alpha_1 < \sum_{i=2}^{k} \alpha_i$. We construct a profile in a series of steps. The number of agents depends on the scoring vector. The linear order $>_1$ of agent 1 is $r_1 >_1 r_2 >_1 \cdots >_1 r_m$. First we give the *manipulative* linear order of agent 2. Then we show how to swap two resources in the manipulative linear order such that a sure resource for agent 2 becomes contested in the original linear order. Agent 2's manipulative preference is constructed as follows: We transform the linear order $>_1$ by a series of swaps. The goal is to maximize the number of sure resources that agent 2 receives. Therefore, all resources in the first bin of agent 1 are swapped with resources in the second bin, the third bin, and so on until all resources of the first bin are in a worse bin. Because of $\alpha_1 < \sum_{i=2}^{k} \alpha_i$ that is possible. Now we consider the second bin of the transformed linear order. Again, we improve the positions of resources in worse bins by swapping them with resources in the second bin. This time, however, we have to pay attention to the fact that we do not improve the position of resources that were sent to worse bins in previous iterations. We continue this procedure until we reach the last bin or no resources are left that can be made sure for agent 2.

More formally, set $>^{(1)} = >_1$. Starting with $j = 1$, swap the resources of $A_j(>^{(j)})$ in decreasing order according to $>^{(1)}$ with those of $A_{j+1}(>^{(j)}) \setminus \bigcup_{l<j} A_l(>^{(1)})$, $A_{j+2}(>^{(j)}) \setminus \bigcup_{l<j} A_l(>^{(1)})$, and so on, until no longer possible. Call the transformed linear order $>^{(j+1)}$. Continue until the transformed linear order $>^{(k)} = >'_2$ is obtained.

In order to transform the manipulative linear order $>'_2$ to the original linear order $>_2$, we look at the rightmost bin $A_w(>'_2)$ which contains a resource of $A_1(>_1)$. We call $A_w(>'_2)$ the worst bin. Because $\alpha_1 < \sum_{i=2}^{k} \alpha_i$ holds, every resource of $A_1(>_1)$ is swapped. Hence, the worst bin is at least the second bin. However, it cannot be the second bin. If it were, then $\alpha_1 \leq \alpha_2$ would hold because otherwise a resource from the first bin would be swapped to at least the third bin. But if all resources of $A_1(>_1)$ are swapped with resources of $A_2(>_2)$, then the next iteration will swap resources that are originally from the first bin with resources of the third bin. This is true because there are at least three different values and swapping is performed in decreasing order according to $>^{(1)}$. Therefore, the worst bin of agent 2 with the manipulative linear order is at least the third bin.

Looking at the resources that agent 1 has in the bin that has the same index as agent 2's worst bin, that is, in $A_w(>_1)$, there is always a resource that is sure for agent 2 but not in the first bin of agent 2 in the manipulation. Note that here we use the fact that the worst bin has to be at least the third bin; otherwise, every resource that is sure for agent 2 would have to be in the first bin. Now we show that such a resource always exists.

Pick an $a \in A_w(>'_2) \cap A_1(>_1)$. There are two possibilities how resource a came into $A_w(>'_2)$:

1. Series of swaps: Resource a was swapped with b which is now sure for agent 2. Because there was at least one preceding swap involving resource a, resource b cannot be in $A_1(>'_2)$.
2. Direct swap: After the first iteration all resources from the bins that are to the left of the worst bin (except for the first bin) were swapped with the resources in the first bin of agent 1. The worst bin has not been completely filled by new resources in the first iteration: If all resources in $A_w(>'_2)$ are there due to a direct swap (that is, in the first iteration), then there is a bin to the right of the worst bin (otherwise, we would have $\alpha_1 = \sum_{i=2}^{k} \alpha_i$, contradicting the case assumption). Then the second iteration swaps a resource from the bin that is to the right of the worst bin with a resource in the second bin. The second bin just contains resources from $A_1(>_1)$. Thus $A_w(>'_2)$ is not the worst bin.

   If there is a resource in $A_w(>'_2)$ that was not swapped in the first iteration, the second iteration picks *such* a resource in order to swap it with a resource in the second bin, which just contains resources from $A_1(>_1)$.

Thus there is a resource $r \in S_2(P')$ with $r \notin A_1(>'_2)$ and $r \in A_w(>_1)$. Now swap r and a in $>'_2$. Resource a remains unavailable to agent 2 because resource r was not in the first bin of agent 2 in the manipulation. However, r becomes contested for agent 2, the first agent having r in the same bin. Call this linear order $>_2$.

Agent 2 may still receive contested resources. By adding additional agents, we prevent that from happening. For each contested resource r_c that agent 2 receives in P', add to profiles P and P' a linear order $>_{r_c}$ which results from swapping resource r_c with resource r_1 in $>_1$. This is always possible because $\alpha_1 < \sum_{i=2}^{k} \alpha_i$ implies that agent 2's contested resources in P' are not in $A_1(>_1)$. Hence, we have a manipulation with $F_s(P')_2 = \{S_2(P')\}$ and $F_s(P)_2 = \{S_2(P') \setminus \{r\}, S_2(P')\}$. This completes Case 2.

In both cases, we have $F_s(P')_2 \succ_2^K F_s(P)_2$. ❑

Example 9 *Let $s = (7,7,7,5,5,5,5,5,3)$ be a given scoring vector (which corresponds to the second case of the proof of Proposition 8). Then the linear order $>_1 = >^{(1)}$ of agent 1 is $r_1\, r_2\, r_3 \,|\, r_4\, r_5\, r_6\, r_7\, r_8 \,|\, r_9$.*

The manipulative linear order $>'_2$ is constructed first:

$$
\begin{aligned}
>^{(2)} &: \quad r_4\, r_5\, r_6 \,|\, r_1\, r_2\, r_3\, r_7\, r_8 \,|\, r_9 \\
>'_2 \;=\; >^{(3)} &: \quad r_4\, r_5\, r_6 \,|\, r_9\, r_2\, r_3\, r_7\, r_8 \,|\, r_1
\end{aligned}
$$

Set profile $I' = (>_1, >'_2)$. Then the set of sure resources for agent 2 in I' is $S_2(I') = \{r_4, r_5, r_6, r_9\}$, the set of contested resources for agent 2 is $C_2(I') = \{r_7, r_8\}$. Now we construct the original linear order $>_2$ of agent 2. The worst bin is $A_3(>'_2)$. A resource in the intersection $A_3(>'_2) \cap A_1(>_1)$ is r_1. The swap that brought r_1 to the third bin was with r_9. Swapping r_1 with r_9 gives $>_2$. Set profile $I = (>_1, >_2)$. Then the set of sure resources for agent 2 in the profile I is $S_2(I) = \{r_4, r_5, r_6\}$. For each contested resource in $C_2(I')$, an agent is added with the linear order $>_1$ except for a swap of this contested resource with r_1. The complete profile P' is:

$$
\begin{aligned}
>_1 &: \quad r_1\, r_2\, r_3 \,|\, r_4\, r_5\, r_6\, r_7\, r_8 \,|\, r_9 \\
>'_2 &: \quad r_4\, r_5\, r_6 \,|\, r_9\, r_2\, r_3\, r_7\, r_8 \,|\, r_1 \\
>_{r_7} &: \quad r_7\, r_2\, r_3 \,|\, r_4\, r_5\, r_6\, r_1\, r_8 \,|\, r_9 \\
>_{r_8} &: \quad r_8\, r_2\, r_3 \,|\, r_4\, r_5\, r_6\, r_7\, r_1 \,|\, r_9
\end{aligned}
$$

and the winning set for agent 2 in profile P' is $F_s(P')_2 = \{S_2(I')\}$. The profile P is the same as P' except for exchanging $>'_2$ with $>_2$. This gives the winning set $F_s(P)_2 = \{S_2(I), S_2(I) \cup \{r_9\}\}$. Overall, the manipulation follows via $F_s(P')_2 \succ_2^K F_s(P)_2$.

Finally, we consider the case of scoring vectors with at least three different values and $\alpha_1 > \sum_{i=2}^{k} \alpha_i$.

Lemma 10 *Let s be a scoring vector with k different values. For each profile P with at least two agents and for each agent j, $\|R_j(P)\| \geq \alpha_1$ and $\|S_j(P)\| \leq \sum_{i=2}^{k} \alpha_i$.*

Proof. For the first inequality, notice that $A_1(>_j) \subseteq R_j(P)$. For the second inequality, no resources that are in the first bin for some agent can be sure for another agent. The bound follows. ❑

Lemma 11 *Let s be a scoring vector with k different values and R be a set of resources. For each agent j, for each responsive linear order $\succ_j$ on 2^R extending $>_j$, and for each set X with $\|X\| \leq \alpha_1$ and $X \neq A_1(>_j)$, we have $A_1(>_j) \succ_j X$.*

Proof. If $X \subset A_1(>_j)$ holds, then the relation follows by responsiveness. Otherwise, consider the sets $A_1(>_j) \setminus X$ and $X \setminus A_1(>_j)$. Let $c = \|A_1(>_j) \cap X\|$ be the number of resources in the intersection. We have $\|A_1(>_j) \setminus X\| = \alpha_1 - c \geq \|X \setminus A_1(>_j)\|$. Hence, each resource in $X \setminus A_1(>_j)$ can be matched to an arbitrary resource in $A_1(>_j) \setminus X$, because for each $y \in A_1(>_j)$ and $x \in X \setminus A_1(>_j)$, we have $y >_j x$. The relation follows by responsiveness. ❑

If the first bin is large, the intuition of the following result on Gärdenfors-strategy-proofness is that the set of sure resources is small (Lemma 10), whereas the (large) first bin is preferred to every small bundle (Lemma 11).

Proposition 12 *For at least two agents and each scoring vector s with $k \geq 3$ different values satisfying $\alpha_1 > \sum_{i=2}^{k} \alpha_i$, F_s is Gärdenfors-strategy-proof.*

Proof. Let R be the set of resources, P be a profile, P' be the same profile as P except for a change of the manipulator's preference, let $>_n$ be the manipulator's true linear order on R, and $\succ_n$ be the manipulator's responsive linear order on 2^R extending $>_n$. Consider the three relations of $F_s(P)_n$ and $F_s(P')_n$ according to Gärdenfors. For each case we exhibit two sets that satisfy the membership criteria in the definition of a Gärdenfors extension and that have the desired relation between them.

Case 1: $F_s(P)_n \subset F_s(P')_n$. In this case we find sets $\pi \in F_s(P)_n$ and $\pi' \in F_s(P')_n \setminus F_s(P)_n$ such that $\pi \succ_n \pi'$ holds. Consider the following subcases with respect to the relationship between $R_n(P)$ and $R_n(P')$:

- **Case 1(a):** $R_n(P') \subset R_n(P)$. Since $F_s(P)_n \subset F_s(P')_n$ implies $R_n(P) \subseteq R_n(P')$, this case never occurs.
- **Case 1(b):** $R_n(P) \subset R_n(P')$. The idea is that adding a resource which is in $R_n(P')$ but not in $R_n(P)$ to the set of sure resources in P' is sufficient for the set not to be in $F_s(P)_n$.

  Suppose that $r \in S_n(P')$ for all $r \in R_n(P') \setminus R_n(P)$. Then $S_n(P') \notin F_s(P)_n$ is true and, by bounding the size of $S_n(P')$ with Lemma 10 and then applying Lemma 11, we have $A_1(>_n) \cup S_n(P) \succ_n S_n(P')$. Now suppose that there exists $r' \in R_n(P') \setminus R_n(P)$ with $r' \notin S_n(P')$. Then we have $S_n(P') \cup \{r'\} \notin F_s(P)_n$ and, by Lemmas 10 and 11 and the fact that $r' \notin A_1(>_n)$, we have $A_1(>_n) \cup S_n(P) \succ_n S_n(P') \cup \{r'\}$.
- **Case 1(c):** $R_n(P) = R_n(P')$. Let $\pi' \in F_s(P')_n \setminus F_s(P)_n$. Then we have $\pi' \neq R_n(P')$ because $R_n(P) \in F_s(P)_n$. Hence, $\pi' \subset R_n(P')$ is true, which implies $R_n(P) \succ_n \pi'$.
- **Case 1(d):** $R_n(P) \not\subset R_n(P')$ and $R_n(P') \not\subset R_n(P)$. This case never occurs, as $F_s(P)_n \subset F_s(P')_n$ implies $R_n(P) \subseteq R_n(P')$.

Case 2: $F_s(P')_n \subset F_s(P)_n$. Then we find sets $\pi \in F_s(P)_n \setminus F_s(P')_n$ and $\pi' \in F_s(P')_n$ such that $\pi \succ_n \pi'$.

- **Case 2(a):** $R_n(P') \subset R_n(P)$. We have $R_n(P) \notin F_s(P')_n$ (otherwise, $R_n(P) \subseteq R_n(P')$). Hence, $R_n(P) \succ_n R_n(P')$ is enough in this case.
- **Case 2(b):** $R_n(P) \subset R_n(P')$. Analogously to Case 1(a).
- **Case 2(c):** $R_n(P) = R_n(P')$. The premises $R_n(P) = R_n(P')$ and $F_s(P')_n \subset F_s(P)_n$ imply that there exists a resource $r \in R_n(P)$ that is contested in P but sure in P'. It follows that $r \notin A_1(>_n)$ is true (otherwise, r would not be sure in P' and contested in P). Thus we have $r \notin A_1(>_n) \cup S_n(P)$, and $A_1(>_n) \cup S_n(P) \notin F_s(P')_n$ holds because r is sure in P'. By Lemmas 10 and 11, $A_1(>_n) \cup S_n(P) \succ_n S_n(P')$.
- **Case 2(d):** $R_n(P) \not\subset R_n(P')$ and $R_n(P') \not\subset R_n(P)$. Analogously to Case 1(d).

Case 3: $F_s(P)_n \not\subset F_s(P')_n$ and $F_s(P')_n \not\subset F_s(P)_n$. In this case we find sets $\pi \in F_s(P)_n \setminus F_s(P')_n$ and $\pi' \in F_s(P')_n \setminus F_s(P)_n$ such that $\pi \succ_n \pi'$ holds.

Case 3(a): $R_n(P') \subset R_n(P)$. We have $R_n(P) \notin F_s(P')_n$ (otherwise, $R_n(P) \subseteq R_n(P')$). Pick a set $\pi' \in F_s(P')_n \setminus F_s(P)_n$. Then $\pi' \subseteq R_n(P')$ is true. This implies that $R_n(P) \succ_n \pi'$.

Case 3(b): $R_n(P) \subset R_n(P')$. Suppose that $A_1(>_n) \cup S_n(P) \notin F_s(P')_n$ holds. Then we can argue similarly to Case 1(b).

Now we show that $A_1(>_n) \cup S_n(P) \in F_s(P')_n$ never holds: We have $S_n(P') \subset A_1(>_n) \cup S_n(P)$. However, if $r \in S_n(P')$ is in $A_1(>_n)$, then $r \in S_n(P)$ holds as well. Hence $S_n(P') \subseteq S_n(P)$ is true. Together with $R_n(P) \subset R_n(P')$ this implies that $F_s(P)_n \subset F_s(P')_n$ because each resource $c \in (S_n(P) \setminus S_n(P')) \cup (R_n(P') \setminus R_n(P))$ is contested in P' (contradiction).

Case 3(c): $R_n(P) = R_n(P')$. $R_n(P) = R_n(P')$ together with $F_s(P)_n \not\subset F_s(P')_n$ implies that there is a resource that is contested in P but sure in P'. Similarly, there is a resource that is contested in P' but sure in P. Overall, we have $S_n(P') \not\subset S_n(P)$ and $S_n(P) \not\subset S_n(P')$. Pick a resource $r \in S_n(P') \setminus S_n(P)$. The remaining argument is analogous to Case 2(c).

Case 3(d): $R_n(P) \not\subset R_n(P')$ and $R_n(P') \not\subset R_n(P)$. We have $R_n(P) \notin F_s(P')_n$ (otherwise, $R_n(P) \subseteq R_n(P')$) and argue similarly to Case 1(b).

This completes the proof. ❑

Example 13 *Let $s = (8,8,8,8,8,5,5,4,2)$ be a scoring vector and $R = \{r_1,\dots,r_9\}$ be the set of resources. Let P be the following profile of three agents:*

$$
\begin{aligned}
>_1 : &\quad r_1\, r_2\, r_3\, r_4\, r_5 \,|\, r_6\, r_7 \,|\, r_8 \,|\, r_9 \\
>_2 : &\quad r_1\, r_3\, r_4\, r_5\, r_7 \,|\, r_2\, r_6 \,|\, r_8 \,|\, r_9 \\
>_3 : &\quad r_1\, r_2\, r_3\, r_4\, r_6 \,|\, r_7\, r_9 \,|\, r_8 \,|\, r_5
\end{aligned}
$$

The winning set for agent 3 consists of $S_3(P) = \{r_6, r_9\}$ and $C_3(P) = \{r_1, r_2, r_3, r_4, r_8\}$. Now consider the following misreported linear order of agent 3:

$$>'_3: \quad r_1\, r_2\, r_3\, r_4\, r_7 \,|\, r_6\, r_9 \,|\, r_8 \,|\, r_5$$

If P' is profile P with linear order $>_3$ replaced by $>'_3$, the new winning set of agent 3 consists of $S_3(P') = \{r_9\}$ and $C_3(P') = \{r_1, r_2, r_3, r_4, r_6, r_7, r_8\}$. Then we have $F_s(P)_3 \subset F_s(P')_3$ and $R_3(P) \subset R_3(P')$ (Case 1(b) of the proof of Proposition 12). We have $R_3(P') \setminus R_3(P) = \{r_7\}$. Hence, $S_3(P') \cup \{r_7\} = \{r_7, r_9\} \in F_s(P')_3$ and $\{r_7, r_9\} \notin F_s(P)_3$. It follows that $A_1(>_3) \cup S_3(P) = \{r_1, r_2, r_3, r_4, r_6, r_9\} \succ_3 \{r_7, r_9\}$ holds because of, e.g., $r_1 >_3 r_7$.

From Propositions 7, 8, and 12, we have the main result:

Theorem 14 *Let s be a scoring vector with k different values. The following three statements are equivalent: (1) F_s is Kelly-strategy-proof; (2) F_s is Gärdenfors-strategy-proof; (3) $k \leq 2$ or $\alpha_1 > \sum_{i=2}^{k} \alpha_i$.*

Corollary 15 *For $s \in \{$borda, lex, ε-qi$\}$, the following three statements are equivalent: (1) F_s is Kelly-strategy-proof; (2) F_s is Gärdenfors-strategy-proof; (3) there are no more than two resources.*

Corollary 16 *$F_{k\text{-app}}$ is Kelly- and Gärdenfors-strategy-proof.*

4 Related Work

Bouveret et al. [2016] and Lang and Rothe [2015] survey the work on fair division of indivisible goods, and Nguyen et al. [2013] that on approximability of social welfare optimization (see also the work of Nguyen et al. [2014]) for cardinal preferences. Since the model that we study is inspired by scoring rules from voting, the chapter by Brams and Fishburn [2002] might serve as an introduction.

Brams and King [2005] assume ordinal preferences over single resources. They look at properties such as maximin, Borda maximin, and envy. Similarly, Brams et al. [2003] assume that each agent has an additively separable preference over the resources, with no side payments allowed. Optimal allocations maximize utilitarian or egalitarian social welfare, where the utility of an agent is the Borda score of the received bundle. Baumeister et al. [2014] extend this model by considering arbitrary scoring vectors and collective utility functions and study scoring allocation rules with respect to axiomatic properties. Strategy-proofness for allocation rules has been studied by, e.g., Pápai [2001], Ehlers and Klaus [2003], and Hatfield [2009].

In the context of social choice, the standard approaches for lifting preferences are due to Kelly [1977], Fishburn [1972], and Gärdenfors [1976]. The axiomatic study of extensions was started by Kannai and Peleg [1984]. A survey is given by Barberà et al. [2004]. The motivation of Kelly [1977] was to abandon the single-valuedness requirement in the Gibbard–Satterthwaite theorem [Gibbard, 1973, Satterthwaite, 1975]. A generalization of strategy-proofness is given, where the focus is on "clear-cut cases" when there is no knowledge of the final selection process. More recently, Brandt [2015] and Brandt and Brill [2011] establish necessary and sufficient conditions for strategy-proofness of irresolute social choice functions using the preference extensions due to Kelly, Fishburn, and Gärdenfors.

5 Conclusions

We have studied resource allocation in a model where resources are indivisible, nonshareable, and of single-unit type. Agents reveal ordinal preferences over single resources only, thus crucially alleviating the elicitation burden. Allocations are complete assignments of all resources to the agents. Winner determination is facilitated through scoring vectors: The ordinal preferences are transformed into utility functions as if the agents' preferences were additively separable. Using these surrogate utility functions, allocations can be assessed with tools of resource allocation with cardinal preferences such as utilitarian social welfare.

We have studied the manipulation problem in this model. This is the question of whether agents can benefit from misreporting their preferences. Specifically, we have taken the point of view of a manipulator where we know the manipulator's responsive preferences over bundles of resources instead of single resources. As multiple allocations can maximize utilitarian social welfare, the manipulator's preferences are lifted to sets of bundles of resources through Kelly and Gärdenfors extensions. We have shown that the structure of the chosen scoring vector characterizes whether the utilitarian allocation correspondence is manipulable or strategy-proof: If there are at most two different values in the scoring vector, then the utilitarian allocation correspondence is Gärdenfors/Kelly-strategy-proof; if there are at least three different values in the scoring vector, then it is Gärdenfors/Kelly-strategy-proof if and only if the number of occurrences of the greatest value in the scoring vector is larger than half the number of goods.

Intuitively speaking, the bottom line is that a utilitarian allocation correspondence becomes strategy-proof if there is high discrepancy between the choice of the scoring vector and *strict* preferences, that is, if the scoring vector assigns the same value to a lot of ranks although agents have strict preferences between resources. This is consistent with the intuition that an allocation correspondence is the closer to strategy-proofness the more oblivious it is to the agents' preferences.

As directions of future research we propose to investigate the complexity of winner determination in this model. Furthermore, a characterization of strategy-proofness for other social welfare measures (e.g., egalitarian social welfare) is still open. Manipulability of allocation rules as compositions of allocation correspondences with tie-breaking mechanisms would constitute an interesting research direction as well.

Acknowledgments

An earlier version of this paper appeared in the Proceedings of IJCAI-2015. We thank the anonymous reviewers for helpful comments and Jérôme Lang for interesting discussions and feedback on an earlier draft of this paper. This work was supported in part by the NRW Ministry for Innovation, Science, and Research and DFG grants RO 1202/14-1, RO 1202/14-2, and RO 1202/15-1.

References

S. Barberà, W. Bossert, and P. Pattanaik. Ranking sets of objects. In S. Barberà, P. Hammond, and C. Seidl, editors, *Handbook of Utility Theory*, volume 2: Extensions, pages 893–977. Kluwer Academic Publisher, 2004.

D. Baumeister, S. Bouveret, J. Lang, N. Nguyen, T. Nguyen, and J. Rothe. Scoring rules for the allocation of indivisible goods. In *Proceedings of the 21st European Conference on Artificial Intelligence*, pages 75–80. IOS Press, August 2014.

S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 73–78. AAAI Press/IJCAI, July 2011.

S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press, 2016. To appear.

S. Brams and P. Fishburn. Voting procedures. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, pages 173–236. North-Holland, 2002.

S. Brams and D. King. Efficient fair division: Help the worst off or avoid envy? *Rationality and Society*, 17(4):387–421, 2005.

S. Brams, P. Edelman, and P. Fishburn. Fair division of indivisible items. *Theory and Decision*, 55 (2):147–180, 2003.

F. Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45(4): 793–804, 2015.

F. Brandt and M. Brill. Necessary and sufficient conditions for the strategyproofness of irresolute social choice functions. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 136–142. ACM Press, July 2011.

I. Caragiannis and A. Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9–10):1655–1671, 2011.

L. Ehlers and B. Klaus. Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. *Social Choice and Welfare*, 21(2):265–280, 2003.

P. Fishburn. Even-chance lotteries in social choice theory. *Theory and Decision*, 3(1):18–40, 1972.

P. Gärdenfors. Manipulation of social choice functions. *Journal of Economic Theory*, 13(2):217–228, 1976.

A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

J.W. Hatfield. Strategy-proof, efficient, and nonbossy quota allocations. *Social Choice and Welfare*, 33(3):505–515, 2009.

E. Hemaspaandra and L. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.

Y. Kannai and B. Peleg. A note on the extension of an order on a set to the power set. *Journal of Economic Theory*, 32(1):172–175, 1984.

J. Kelly. Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, 45(2):439–446, 1977.

J. Lang and J. Rothe. Fair division of indivisible goods. In J. Rothe, editor, *Economics and Computation. An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, chapter 8, pages 493–550. Springer-Verlag, 2015.

N. Nguyen, T. Nguyen, M. Roos, and J. Rothe. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Journal of Autonomous Agents and Multi-Agent Systems*, 28(2):256–289, 2014.

T. Nguyen, M. Roos, and J. Rothe. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence*, 68(1–3):65–90, 2013.

S. Pápai. Strategyproof and nonbossy multiple assignments. *Journal of Public Economic Theory*, 3 (3):257–271, 2001.

M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2): 187–217, 1975.

Nhan-Tam Nguyen, Dorothea Baumeister, and Jörg Rothe
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany
Email: `{nguyen,baumeister,rothe}@cs.uni-duesseldorf.de`