ECAI 2020
G.D. Giacomo et al. (Eds.)

© 2020 The authors and IOS Press.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/FAIA200076

Preferences over Rankings and How to Control Them Using Rewards

Marco Faella[1] and **Luigi Sauro**[2]

Abstract. We study the rational preferences of agents participating in a mechanism whose outcome is a weak order among participants. We propose a set of self-interest axioms and characterize the mutual relationships between all subsets thereof. We then assume that the mechanism can assign monetary rewards to the agents, in a way that is consistent with the weak order. We show that the mechanism can induce specific classes of preferences by suitably choosing the assigned rewards, even in the absence of tie breaking.

1 Introduction

In several applications, a set of agents competes in a mechanism whose output is a ranking among the agents themselves. Tournaments [7], recruiting competitions, some types of auctions [12] and elections [4, 8, 6], or partitioning students in groups with homogeneous level of ability [10] can all be considered in this way. In order to design better mechanisms of this sort, we need tools to predict the most likely behavior of the agents. If the agents are rational and know the rules of the mechanism, they will project their preference among different outcomes back to their choice of actions. That is indeed the basic premise of game theory. Therefore, the first step to predict and control the agents' behavior is to analyze the ordinal preference that they are likely to exhibit over the outcomes of the mechanism.

If the outcomes are linear orders among agents, there is virtually no doubt about the way in which an agent compares two rankings: the higher their placement in the ranking, the better.[3] On the other hand, a ranking mechanism can in principle fail to distinguish two agents [1, 14, 11], ending up with one or more ties (formally, it returns a weak order). In that case, a tie-breaking phase might be in order, usually based either on randomization [9] or on a fixed order between competing entities [13]. However, when applied to agents, these domain-independent tie-breaking techniques may appear arbitrary or, even worse, prejudicial.

In this paper, we investigate the space of preferences that agents may exhibit if the output of the mechanism is a weak order, and then study a family of fair alternatives to tie-breaking, that promote competition while treating all agents equally and avoiding randomization. To this aim, we initially assume that an arbitrary mechanism produces a weak order among the participating agents. It is immediate to see that in that case the preference of a rational agent is far from uniquely specified, as ties can be perceived as more or less harmful depending on the context. To put some order in the space of preferences, we introduce some natural conditions reflecting different aspects of the rational behaviour of a self-interested agent and study the theories deriving from choosing a generic subset of those conditions. We show that only a relatively small number of theories are distinct and we characterize the taxonomy of their mutual relationships.

We then investigate how specific classes of preferences may be induced by assigning *quantitative rewards* to the participating agents. Specifically, we consider the case where the mechanism applies a reward policy to convert a weak order among agents to a quantitative reward for each participant. Adding a reward policy resolves the nondeterminism concerning the agents' preferences: by assuming that each agent is interested in maximising its own reward, we obtain a well-specified preference for each agent.

As a first observation, reward policies provide a stronger justification to the aforementioned rationality conditions: it is sufficient to impose very general fairness assumptions on the policy to induce preferences that satisfy at least the weakest rationality conditions. Notice that this approach is independent of the specific ranking-producing mechanism, which is instead the focus of standard mechanism design. As a consequence, our reward policies can be applied to a wide range of concrete mechanisms.

We then focus on reward policies that split a fixed jackpot over the agents, in a way that is consistent with their placements in the output ranking. Notably, in this setting some preference theories collapse into a single one. Finally, even if reward policies can in principle perform tie-breaking, we characterize a family of policies that do not perform tie-breaking and yet promote competition by inducing agents to overcome each other. For example, in the absence of a reward mechanism the agents may adopt a preference that is insensitive to ties. Such a preference does not induce agents to act in a way that pushes same-level competitors down in the ranking. In turn, this lack of incentive may compromise the success of the competition from the point of view of the organizer. We show that a suitable reward policy can prevent this problem by ensuring that the agents prefer configurations where as many competitors as possible lie strictly below them and as few as possible lie strictly above them.

The paper is structured as follows. In Section 2 we introduce preliminary definitions as well as motivating examples of different preferences an agent may adopt. In Section 3 we propose a list of self-interest conditions over preferences, modelling different aspects of rationality and we study their properties and mutual relationships. In Section 4 we combine the aforementioned conditions and provide a complete hierarchy of the resulting theories. Section 5 concerns reward policies that can be applied to rankings. A short section with concluding remarks ends the paper.

2 Preliminaries

Recall the following definitions: a *pre-order* is a reflexive and transitive binary relation, a *weak order* is a total pre-order, and a *linear*

[1] University of Naples Federico II, Italy, email: m.faella@unina.it
[2] University of Naples Federico II, Italy, email: luigi.sauro@unina.it
[3] Clearly, we are implicitly assuming that the output ranking represents some kind of value judgement over the agents.

order is an antisymmetric weak order.

Given a pre-order $\sqsubseteq$, we denote its asymmetric part by $\sqsubset$, its symmetric part by $\equiv$, and its symmetric complement by $\not\lessgtr$. Formally, we have:

$$a \sqsubset b \text{ iff } a \sqsubseteq b \text{ and } b \not\sqsubseteq a$$
$$a \equiv b \text{ iff } a \sqsubseteq b \text{ and } b \sqsubseteq a$$
$$a \not\lessgtr b \text{ iff } a \not\sqsubseteq b \text{ and } b \not\sqsubseteq a\,.$$

We say that a pre-order $\sqsubseteq_2$ *perfects* another pre-order $\sqsubseteq_1$ if $a \sqsubset_1 b$ implies $a \sqsubset_2 b$, and $a \not\lessgtr_2 b$ implies $a \not\lessgtr_1 b$. Intuitively, $\sqsubseteq_2$ exhibits at least the same strong preferences of $\sqsubseteq_1$, but it may distinguish items that are equivalent for $\sqsubseteq_1$, as well as distinguishing or equating items that are incomparable for $\sqsubseteq_1$. This notion is incomparable to the classical notion of refinement (that is, inclusion) between relations. Perfectioning is itself a pre-order, whose bottom element is the identity relation and whose top elements are the linear orders.

We assume that an undescribed process, such as a competition or a vote, produces a weak order on a finite set of agents A. When a weak order involves agents we call it a *ranking* and we denote by $\mathcal{R}(A)$ the set of all possible rankings over A. Unless differently specified, we consider a fixed set A containing at least two agents. In the following sections we will investigate how the involved agents compare different rankings and establish a preference over them. For a ranking $\sqsubseteq$ and an agent a, let $bl_a(\sqsubseteq)$, $eq_a(\sqsubseteq)$, $ab_a(\sqsubseteq)$ be the partition of A into the agents that are strictly below, equivalent, and strictly above a, respectively. Moreover, let $lvl_a^\top(\sqsubseteq)$ be the length of the longest $\sqsubset$-chain starting from a, which can be recursively defined as follows:

$$lvl_a^\top(\sqsubseteq) = \begin{cases} 1 & \text{if } ab_a(\sqsubseteq) = \emptyset \\ \max_{b \in ab_a(\sqsubseteq)} lvl_b^\top(\sqsubseteq) + 1 & \text{otherwise.} \end{cases}$$

One can dually define $lvl_a^\bot(\sqsubseteq)$ as the length of the longest $\sqsubset$-chain *ending* in a. Informally, this measures the level of a, starting from the bottom.

Given a ranking $\sqsubseteq$, the set of agents A can be partitioned in those agents a that are on the top level ($lvl_a^\top(\sqsubseteq) = 1$), those agents at level two and so on. Notice that a ranking $\sqsubseteq_2$ prefecting another ranking $\sqsubseteq_1$ means that $\sqsubseteq_2$ can split some levels from $\sqsubseteq_1$ into multiple levels.

Finally, a *preference relation for an agent* a (in short, a-preference) is a weak order on $\mathcal{R}(A)$.

2.1 A Catalog of Preferences

In this section we show a short catalog of possible ways to compare two rankings from the point of view of a participant.

To identify a preference, we use a uniform naming system based on the following abbreviations:

Abbreviation	Meaning
$\top$	$lvl_a^\top(\sqsubseteq)$
$\bot$	$-lvl_a^\bot(\sqsubseteq)$
ab	$\lvert ab_a(\sqsubseteq)\rvert$
eq	$\lvert eq_a(\sqsubseteq)\rvert$
bl	$\lvert bl_a(\sqsubseteq)\rvert$

We use the above abbreviations as a superscript to indicate that a preference tries to minimise the corresponding quantity. For example, P_a^{ab} is the a-preference that minimises $|ab_a(\sqsubseteq)|$, i.e., the number of agents strictly above a in the ranking. Formally, $P_a^{\mathrm{ab}}(\sqsubseteq_1, \sqsubseteq_2)$ holds iff $|ab_a(\sqsubseteq_1)| \geq |ab_a(\sqsubseteq_2)|$. We also write $P_a^{x,y}$ for the preference that tries to minimise first quantity x and then quantity y, lexicographically. Finally, we allow simple arithmetic expressions, like $P_a^{\mathrm{ab+eq}}$ for the preference that minimises the sum of $|ab_a(\sqsubseteq)|$ and $|eq_a(\sqsubseteq)|$.

Here is a selection of possible preferences based on the above measures:

- Level-based. $P_a^\top$ prefers to minimise the level of a in the ranking, i.e., $lvl_a^\top(\sqsubseteq)$. In particular, it is insensitive to ties.
- Level-based top-k. For a positive integer k, $P_a^{\top_k}$ prefers a to be within the first k levels in the ranking. This preference only distinguishes two classes of rankings: those where a sits in one of the top k levels, and all other rankings. Any ranking in the first class is (strongly) preferred to any ranking in the second class.
- Relative. $P_a^{\mathrm{ab,eq}}$ prefers having fewer agents above a; equal that, it prefers to have fewer agents tied with a. It is equivalent to minimising the vector $\big(|ab_a(\sqsubseteq)|, |eq_a(\sqsubseteq)|\big)$, lexicographically.
- Best linearization. P_a^{bst} judges a ranking the same as its linear extension where a has the best position. The canonical name for this preference is P_a^{ab}, because it is equivalent to minimising $|ab_a(\sqsubseteq)|$.
- Worst linearization. P_a^{wst} is the dual of P_a^{bst}. The canonical name for this preference is $P_a^{\mathrm{ab+eq}}$, because it is equivalent to minimising $|ab_a(\sqsubseteq)| + |eq_a(\sqsubseteq)|$.

The preference $P_a^\top$ may reflect a conscientious participant when a group of students is being partitioned into different levels of ability. Similarly, the preference $P_a^{\top_k}$ can be adopted by the candidates of a multi-winner voting systems where different scoring rules (e.g. plurality score or Borda score) can be used to select those agents that are in the top-k placements.

In [7] the authors use the preference $P_a^{\mathrm{ab,eq}}$ to model the fact that a player participating in a tournament generally aims at prevailing over the opponents.

Preferences P_a^{bst} and P_a^{wst} can be adopted whenever an unknown rule is used to resolve ties. In particular, they reflect an optimistic or pessimistic attitude, respectively.

The previous examples show how, depending on the specific context, different ways to compare rankings can be adopted. At the same time, it is evident that not all possible weak orders on $\mathcal{R}(A)$ make sense. For instance, a preference where coming in last position is better than being first is poorly tenable in any context. In the next section we investigate which preferences are plausible from the point of view of a rational and self-interested agent.

3 Preference Axioms

In this section we formalize some *preference axioms* and investigate their mutual relationship. First, we introduce some basic binary relations on rankings. Let $\sqsubseteq_1, \sqsubseteq_2 \in \mathcal{R}(A)$ be two rankings and $a \in A$ an agent:

Same-context. The restrictions of $\sqsubseteq_1$ and $\sqsubseteq_2$ to $(A \setminus \{a\})^2$ coincide. If we need to emphasize the arguments, this property can also be denoted by the extended notation $\mathsf{SC}_a(\sqsubseteq_1, \sqsubseteq_2)$.

Dominance. For all $b \in A$, if $b \sqsubset_1 a$ then $b \sqsubset_2 a$, and if $b \equiv_1 a$ then $b \sqsubseteq_2 a$. Extended notation: $\mathsf{Dom}_a(\sqsubseteq_1, \sqsubseteq_2)$.

Improvement. There exists $b \in A$ such that either $a \equiv_1 b$ and $b \sqsubset_2 a$, or $a \sqsubset_1 b$ and $b \sqsubseteq_2 a$. Extended notation: $\mathsf{Impr}_a(\sqsubseteq_1, \sqsubseteq_2)$.

Swap. There exists $b \in A$ such that $a \sqsubset_1 b$ and $\sqsubseteq_2$ is obtained from $\sqsubseteq_1$ by switching a and b. Extended notation: $\mathsf{Swap}_a(\sqsubseteq_1, \sqsubseteq_2)$.

Notice that, by definition, $\mathsf{Dom}_a(\sqsubseteq_1, \sqsubseteq_2)$ is equivalent to the conjunction of the following conditions:

$$bl_a(\sqsubseteq_1) \subseteq bl_a(\sqsubseteq_2) \quad \text{(I)}$$

$$eq_a(\sqsubseteq_1) \subseteq eq_a(\sqsubseteq_2) \cup bl_a(\sqsubseteq_2) \quad \text{(II)}$$

and also equivalent to the conjunction of the following:

$$ab_a(\sqsubseteq_2) \subseteq ab_a(\sqsubseteq_1) \quad \text{(III)}$$

$$ab_a(\sqsubseteq_2) \cup eq_a(\sqsubseteq_2) \subseteq ab_a(\sqsubseteq_1) \cup eq_a(\sqsubseteq_1)\,. \quad \text{(IV)}$$

Then, assuming that a ranking $\sqsubseteq_2$ dominates $\sqsubseteq_1$ for agent a, the following lemma establishes other equivalences depending on whether $\sqsubseteq_2$ also improves $\sqsubseteq_1$.

Lemma 1 *Let $\sqsubseteq_1, \sqsubseteq_2 \in \mathcal{R}(A)$ be two rankings such that $\mathsf{Dom}_a(\sqsubseteq_1, \sqsubseteq_2)$. The following are equivalent:*

1. *Not $\mathsf{Impr}_a(\sqsubseteq_1, \sqsubseteq_2)$;*
2. *$\mathsf{Dom}_a(\sqsubseteq_2, \sqsubseteq_1)$;*
3. *$\sqsubseteq_1$ and $\sqsubseteq_2$ are* 3-tier equivalent *for a, i.e., $bl_a(\sqsubseteq_1) = bl_a(\sqsubseteq_2)$, $eq_a(\sqsubseteq_1) = eq_a(\sqsubseteq_2)$, and $ab_a(\sqsubseteq_1) = ab_a(\sqsubseteq_2)$.*

Moreover, the following are equivalent:

4. *$\mathsf{Impr}_a(\sqsubseteq_1, \sqsubseteq_2)$;*
5. *not $\mathsf{Dom}_a(\sqsubseteq_2, \sqsubseteq_1)$;*
6. *one of the two inclusions in* (I) *or* (II) *(resp.* (III) *or* (IV)*) is strict.*

Proof: Regarding the first triple of equivalences, the following chain of implications holds: if $\mathsf{Impr}_a(\sqsubseteq_1, \sqsubseteq_2)$ does not hold, the definition of dominance boils down to: for all $b \in A$, if $b \sqsubset_1 a$ then $b \sqsubset_2 a$, and if $b \equiv_1 a$ then $b \equiv_2 a$. This clearly implies that $\sqsubseteq_1$ and $\sqsubseteq_2$ are 3-tier equivalent for a. Then, assume $\sqsubseteq_1$ and $\sqsubseteq_2$ are 3-tier equivalent for a. By conditions (III) and (IV), it immediately holds that also $\sqsubseteq_2$ dominates $\sqsubseteq_1$ for a. Finally, assume that also $\sqsubseteq_2$ dominates $\sqsubseteq_1$ for a. In particular, we have that $b \sqsubset_1 a$ is equivalent to $b \sqsubset_2 a$ and $b \equiv_1 a$ is equivalent to $b \equiv_2 a$. Consequently, none of the two cases in the definition of Improvement is satisfied.

The second triple of equivalences is a direct consequence of the first one and conditions (I) or (II) (resp., (III) or (IV)). □

The properties introduced above allow us to describe four notions that represent a change for the better for agent a:

- a trades place with an agent that is higher up in the order (Swap);
- a moves up in the order, and all other agents stay put (Same-context and Improvement);
- all other agents do not cross the position of a (Dominance).
- a moves up in the order, and all other agents do not cross the position of a (Dominance and Improvement).

In any of these four cases, one may think that a rational self-interested a should prefer the new situation to the old one, at least weakly. In fact, in the following we show that there are reasonable scenarios where some of the above are undesired. For each notion, we propose two axioms prescribing weak and strong preference, respectively. We break the symmetry only for Dom, since Dom is a reflexive relation and hence it cannot imply strong preference (it would be unsatisfiable). To obtain a strong axiom, we pair Dom and Impr, which together essentially define strong dominance (see Lemma 1).

To simplify the notation, hereafter we refer by default to the preference of agent a, and omit the corresponding subscript. Moreover, by $PP(\sqsubseteq_1, \sqsubseteq_2)$ we mean that preference P strictly prefers ranking $\sqsubseteq_2$ to ranking $\sqsubseteq_1$ (that is, $P(\sqsubseteq_1, \sqsubseteq_2)$ holds and $P(\sqsubseteq_2, \sqsubseteq_1)$ does not hold). We obtain the following seven axioms:

$$\mathsf{Swap}(\sqsubseteq_1, \sqsubseteq_2) \longrightarrow P(\sqsubseteq_1, \sqsubseteq_2) \quad \text{(Sw)}$$

$$\mathsf{Swap}(\sqsubseteq_1, \sqsubseteq_2) \longrightarrow PP(\sqsubseteq_1, \sqsubseteq_2) \quad \text{(SSw)}$$

$$\mathsf{SC}(\sqsubseteq_1, \sqsubseteq_2) \wedge \mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2) \longrightarrow P(\sqsubseteq_1, \sqsubseteq_2) \quad \text{(SCI)}$$

$$\mathsf{SC}(\sqsubseteq_1, \sqsubseteq_2) \wedge \mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2) \longrightarrow PP(\sqsubseteq_1, \sqsubseteq_2) \quad \text{(SSCI)}$$

$$\mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2) \longrightarrow P(\sqsubseteq_1, \sqsubseteq_2) \quad \text{(Dom)}$$

$$\mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2) \wedge \mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2) \longrightarrow P(\sqsubseteq_1, \sqsubseteq_2) \quad \text{(DomI)}$$

$$\mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2) \wedge \mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2) \longrightarrow PP(\sqsubseteq_1, \sqsubseteq_2) \quad \text{(SDomI)}$$

A preference P satisfies one of the above axioms if the axiom holds for all rankings $\sqsubseteq_1, \sqsubseteq_2$. Given an axiom α, we denote by $\mathit{Pref}_A(\alpha)$ the set of preferences satisfying α. As customary, we say that the axiom α implies another axiom β, denoted by $\alpha \implies \beta$, if for all sets of agents A and all preferences P, if P satisfies α then it also satisfies β (i.e., $\mathit{Pref}_A(\alpha) \subseteq \mathit{Pref}_A(\beta)$).

The following theorem characterizes the implication relationships holding between the preference axioms.

Theorem 1 *The only implications holding among the preference axioms are displayed in Figure 1 (including the transitive closure of the arrows).*

Proof: The implications $SSw \implies Sw$, $Dom \implies DomI$, $SDomI \implies DomI$, and $SSCI \implies SCI$ are obvious by definition. The implications between the other axioms derive from the following two implications between their premises.

Implication 1:

$$\mathsf{Swap}(\sqsubseteq_1, \sqsubseteq_2) \text{ implies } \mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2) \wedge \mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2)\,.$$

Assume $\mathsf{Swap}(\sqsubseteq_1, \sqsubseteq_2)$ and let b be the agent that trades place with a. Notice that b is strictly above a in $\sqsubseteq_1$. To prove that Dominance holds, let c be such that $c \sqsubset_1 a$. Since a moves up in the order, we have that $c \sqsubset_2 a$. If instead we consider c s.t. $c \equiv_1 a$, for the same reason we have that $c \sqsubseteq_2 a$. To prove that $\mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2)$ holds, simply use b as witness. This claim proves that $DomI \implies Sw$ and $SDomI \implies SSw$.

Implication 2:

$$\mathsf{SC}(\sqsubseteq_1, \sqsubseteq_2) \wedge \mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2) \text{ implies } \mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2) \wedge \mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2)\,.$$

It suffices to prove that $\mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2)$ holds. Let $\sqsubseteq_1, \sqsubseteq_2$ be two rankings satisfying SC and Impr, and let b be the witness agent for Impr. To prove dominance, let c be an agent s.t. $c \sqsubset_1 a$ (if there is no such agent, this part of the claim is vacuously true). By definition of $\mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2)$, $a \sqsubseteq_1 b$ and therefore $b \neq c$. It follows that $c \sqsubset_1 b$ and, by $\mathsf{SC}(\sqsubseteq_1, \sqsubseteq_2)$, $c \sqsubset_2 b$. Since $b \sqsubset_2 a$, by transitivity we obtain $c \sqsubset_2 a$, as required by dominance.

Next, let c be s.t. $c \equiv_1 a$. If $b = c$ then $c \sqsubseteq_2 a$, as required by dominance. Otherwise, we distinguish two further cases, according to the disjunction in the definition of improvement. If $a \equiv_1 b$ then $c \equiv_1 b$ and, by $\mathsf{SC}(\sqsubseteq_1, \sqsubseteq_2)$, $c \equiv_2 b$. By definition of improvement, $b \sqsubset_2 a$ and therefore $c \sqsubset_2 a$, as required by dominance. Finally, if $a \sqsubset_1 b$ then $c \sqsubset_1 b$ and, by the same-context condition, $c \sqsubset_2 b$. Since $b \sqsubseteq_2 a$, by transitivity we obtain $c \sqsubset_2 a$. This claim proves that $DomI \implies SCI$ and that $SDomI \implies SSCI$.

Each non-implication can be proved by a suitable counter-example. Due to space limitations, we omit the details. □

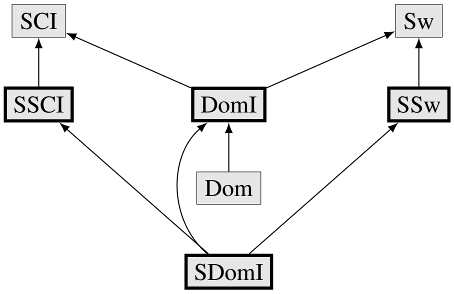

Figure 1. Hasse diagram for implication order between self-interest axioms. Thick boxes denote axioms that prescribe strong preference.

All the preferences in Section 2.1 satisfy a notable property: agent a takes into account only its own placement considering all the other agents in the same way. More formally, we introduce the condition $\mathsf{Ind\_to\_others}(\sqsubseteq_1, \sqsubseteq_2)$, meaning that $\sqsubseteq_2$ is obtained from $\sqsubseteq_1$ by permutating the agents in $A \setminus \{a\}$. Notice that $\mathsf{Ind\_to\_others}(\cdot, \cdot)$ is an equivalence relation. Then, we say that a is *indifferent to the others* if $\mathsf{Ind\_to\_others}(\sqsubseteq_1, \sqsubseteq_2)$ implies that $\sqsubseteq_1$ and $\sqsubseteq_2$ are equivalent for a.

$$\mathsf{Ind\_to\_others}(\sqsubseteq_1, \sqsubseteq_2) \longrightarrow P(\sqsubseteq_1, \sqsubseteq_2) \wedge P(\sqsubseteq_2, \sqsubseteq_1) \qquad \text{(Ind)}$$

The following theorem states that Ind is independent of the other axioms in the sense that Ind does not imply any other axiom and it is not implied by any of them.

Theorem 2 *Axiom Ind is independent of the other axioms.*

Proof: According to the taxonomy proved in Theorem 1, it suffices to prove the following non-implications.

$Ind \not\Longrightarrow SCI$ and $Ind \not\Longrightarrow Sw$. Consider the preference $P^{-\top}$ where agent a prefers to maximize its level. It is straightforward to see that $P^{-\top}$ satisfies Ind and yet it does not satisfy neither SCI nor Sw.

$SDomI \not\Longrightarrow Ind$ and $Dom \not\Longrightarrow Ind$. Pick an agent $b \neq a$ and consider the preference relation denoted by $\hat{P} = P^{\mathrm{ab,eq},a \geq b}$, that perfects $P^{\mathrm{ab,eq}}$ in the following way: when two rankings are equivalent for $P^{\mathrm{ab,eq}}$, $\hat{P}$ strongly prefers the ranking where agent a comes before b.

We first show that $\hat{P}$ satisfies both $SDomI$ and Dom. Assume that $\mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2)$ holds. By (III), we know that $ab(\sqsubseteq_2) \subseteq ab(\sqsubseteq_1)$; if the containment is strict we immediately obtain that agent a strictly prefers $\sqsubseteq_2$. Conversely, if $ab(\sqsubseteq_2) = ab(\sqsubseteq_1)$, condition (IV) implies that $eq(\sqsubseteq_2) \subseteq eq(\sqsubseteq_1)$. Again if $eq(\sqsubseteq_2) \subset eq(\sqsubseteq_1)$, then $\hat{P}$ strictly prefers $\sqsubseteq_2$. Conversely, assume that $eq(\sqsubseteq_1) = eq(\sqsubseteq_2)$. In this case $\sqsubseteq_1$ and $\sqsubseteq_2$ are 3-tier equivalent and hence $b \sqsubseteq_1 a$ iff $b \sqsubseteq_2 a$; consequently $\sqsubseteq_1$ and $\sqsubseteq_2$ are equivalent for a. This concludes the proof that $\hat{P}$ satisfies Dom. Assume now that also $\mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2)$ holds. As before, if $ab(\sqsubseteq_2) \subset ab(\sqsubseteq_1)$, then $\hat{P}$ strictly prefers $\sqsubseteq_2$. When $ab(\sqsubseteq_2) = ab(\sqsubseteq_1)$, instead, Lemma 1 ensures that $eq(\sqsubseteq_2) \subset eq(\sqsubseteq_1)$ and hence also in this case $\hat{P}$ strictly prefers $\sqsubseteq_2$. It remains to show that $\hat{P}$ violates Ind. To this aim, consider the linear orders $c \sqsubset_1 a \sqsubset_1 b$ and $b \sqsubset_2 a \sqsubset_2 c$. Clearly, $\mathsf{Ind\_to\_others}(\sqsubseteq_1, \sqsubseteq_2)$ holds, since $\sqsubseteq_2$ is obtained from $\sqsubseteq_1$ by permutating agents b and c. However, $\hat{P}$ strictly prefers $\sqsubseteq_2$ where a leaves b behind. □

4 Preference Theories

As we have seen in Section 3, there are some reasonable cases where not all the preference axioms can be expected to hold. Nevertheless, it seems equally reasonable to combine multiple axioms to further circumscribe the class of preferences an agent may adopt. In this section we investigate how axioms can be combined to form *preference theories*. We initially focus on the axioms in Figure 1, treating axiom Ind separately.

Let $\mathcal{AX}$ denote the set of axioms in Figure 1. As usual, given a subset $\mathcal{T} \subseteq \mathcal{AX}$, $Pref_A(\mathcal{T})$ is the set of preferences satisfying all the axioms in $\mathcal{T}$. As for single axioms, we say that a theory $\mathcal{T}$ implies $\mathcal{T}'$, denoted by $\mathcal{T} \implies \mathcal{T}'$, if for all sets of agents A it holds $Pref_A(\mathcal{T}) \subseteq Pref_A(\mathcal{T}')$.

We claim that axioms SCI and Sw are primitive and should be satisfied by any preference. This claim is supported by the following observations: *(i)* axioms SCI and Sw are mutually independent and are not subsumed by any other axiom (see Figure 1), and *(ii)* all "plausible" preferences that we have been able to define satisfy those two axioms. Consequently, we will consider only those theories that imply $\{SCI, Sw\}$. The following preliminary lemma shows that DomI and SSCI together imply SSw.

Lemma 2 *If a preference relation satisfies DomI and SSCI, then it satisfies SSw.*

Proof: Let P be a preference satisfying DomI and SSCI, and let $\sqsubseteq_1$ and $\sqsubseteq_2$ be two rankings such that $\mathsf{Swap}(\sqsubseteq_1, \sqsubseteq_2)$ holds. By definition, there is an agent b different from a such that $a \sqsubset_1 b$ and $\sqsubseteq_2$ is obtained from $\sqsubseteq_1$ by switching a and b. Let $\sqsubseteq_3$ be the ranking that coincides with $\sqsubseteq_1$, except for the position of a, which rises to the level of b, so that $a \equiv_3 b$. Notice that it holds $\mathsf{SC}(\sqsubseteq_1, \sqsubseteq_3)$ and $\mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_3)$, the latter thanks to agent b, who is strictly better than a in $\sqsubseteq_1$, and equivalent to a in $\sqsubseteq_3$. Therefore, by SSCI we have that P strictly prefers $\sqsubseteq_3$ to $\sqsubseteq_1$. Now, the only difference between $\sqsubseteq_3$ and $\sqsubseteq_2$ consists in b moving down to the level that a occupies in $\sqsubseteq_1$. As a consequence, it holds $\mathsf{Dom}(\sqsubseteq_3, \sqsubseteq_2)$ and $\mathsf{Impr}(\sqsubseteq_3, \sqsubseteq_2)$. By axiom DomI, P weakly prefers $\sqsubseteq_2$ over $\sqsubseteq_3$. By transitivity, P strongly prefers $\sqsubseteq_2$ over $\sqsubseteq_1$. So, P satisfies SSw. □

Our next result states that the theories in Figure 4 cover all possible combinations of axioms in $\mathcal{AX}$ that imply at least axioms SCI and Sw.

Theorem 3 *For all $\mathcal{T} \subseteq \mathcal{AX}$, if $Pref_A(\mathcal{T}) \subseteq Pref_A(SCI, Sw)$ then $Pref_A(\mathcal{T})$ is equal to one of the canonical theories in Figure 4.*

Proof: Let $\mathcal{T} \subseteq \mathcal{AX}$. Obtain $\mathcal{T}' \subseteq \mathcal{T}$ by removing from $\mathcal{T}$ all axioms that are redundant due to the implications in Figure 1. Clearly, the preference theories of $\mathcal{T}$ and $\mathcal{T}'$ coincide.

If $SDomI \in \mathcal{T}'$, then either $\mathcal{T}' = \{Dom, SDomI\}$ or $\mathcal{T}' = \{SDomI\}$. In both cases, the thesis holds because $\mathcal{T}'$ is one of the the theories in Figure 4. In the rest of the proof we can assume that $SDomI \notin \mathcal{T}'$.

Next, assume that $Dom \in \mathcal{T}'$. Then, $\mathcal{T}'$ is one of the following theories: (i) $\{Dom\}$, (ii) $\{Dom, SSw\}$, (iii) $\{Dom, SSCI\}$, or (iv) $\{Dom, SSw, SSCI\}$. The first three theories are canonical whereas, from Lemma 2 and the fact that Dom implies $DomI$, the theory $\{Dom, SSw, SSCI\}$ is equivalent to $\{Dom, SSCI\}$. In the rest of the proof we can assume that $Dom \notin \mathcal{T}'$.

By a similar argument as for Dom, it can be proved that the only canonical theories that include $DomI$ are $\{DomI\}$,

$\{DomI, SSw\}$, and $\{DomI, SSCI\}$, which are all present in Figure 4. Thereafter, we assume that $DomI \notin \mathcal{T}'$. The remaining theories which include either SSw or $SSCI$ are $\{SSw, SSCI\}$, $\{Sw, SSCI\}$, and $\{SSw, SCI\}$. All of them are canonical. Finally, $\{Sw, SCI\}$ is canonical. □

As the following theorem shows, the hierarchy in Figure 4 among all canonical theories is correct and complete. In particular, Theorem 4 makes use of the pre-order P^{dom} on $\mathcal{R}(A)$ such that $P^{dom}(\sqsubseteq_1, \sqsubseteq_2)$ iff $\mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2)$. Notice that P^{dom} is not a preference because it is not a total order.

Theorem 4 *Let $\mathcal{T}_1, \mathcal{T}_2$ be two of the theories in Figure 4. It holds $\mathcal{T}_1 \implies \mathcal{T}_2$ if and only if there is a path from $\mathcal{T}_1$ to $\mathcal{T}_2$ in Figure 4.*

Proof: It is easy to verify that most implications in Figure 4 directly derive from Theorem 1. For instance, since $SDomI$ implies both $DomI$ and $SSCI$, then we clearly have that $\{SDomI\}$ implies $\{DomI, SSCI\}$. The only two implications that do not directly derive from Theorem 1, namely the fact that $\{DomI, SSCI\}$ implies both $\{DomI, SSw\}$ and $\{SSCI, SSw\}$, follow from Lemma 2 instead. Next, we prove the non-implications that separate all theories in Figure 4.

$\{SDomI\} \not\Longrightarrow \{Dom\}$ and $\{Dom\} \not\Longrightarrow \{Dom, SSw\}$. Also these non-implications directly derive from Theorem 1 where we have already proved that SDomI does not imply Dom, which in turn does not imply SSw (see Figure 1).

$\{SSCI, SSw\} \not\Longrightarrow \{DomI\}$. Consider the preference $P^{\top,\mathrm{eq}}$. Since $P^{\top,\mathrm{eq}}$ perfects $P^{\top}$ it is immediate to see that it satisfies SSw. We show that also SSCI is satisfied. Assume that $\sqsubseteq_1$ and $\sqsubseteq_2$ are two rankings that agree on the relative order of all agents different from a, i.e., the rankings satisfy $\mathsf{SC}(\sqsubseteq_1, \sqsubseteq_2)$. If also $\mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2)$ holds, then two cases may occur. In the first case, $a \sqsubset_1 b$ and $b \sqsubseteq_2 a$ for some b, we have that the level of a from the top is strictly lower in $\sqsubseteq_2$ and hence $P^{\top,\mathrm{eq}}$ strictly prefers $\sqsubseteq_2$ to $\sqsubseteq_1$. In the second case, $a \equiv_1 b$ and $b \sqsubset_2 a$ for some b, a remains in the same level, but the set of other agents that are at the same level as a is strictly smaller in $\sqsubseteq_2$ and hence again $P^{\top,\mathrm{eq}}$ strictly prefers $\sqsubseteq_2$ to $\sqsubseteq_1$. This proves that $P^{\top,\mathrm{eq}}$ satisfies SSCI.
It remains to show that $P^{\top,\mathrm{eq}}$ does not satisfy DomI. Consider four agents, $A = \{a, b, c, d\}$, and a ranking $\sqsubseteq_1$ such that $b \equiv_1 c$, $a \sqsubset_1 b$ and $a \equiv_1 d$. Then, let $\sqsubseteq_2$ be the linear order $d \sqsubset_2 a \sqsubset_2 c \sqsubset_2 b$. It is straightforward to see that both $\mathsf{Dom}(\sqsubseteq_1, \sqsubseteq_2)$ and $\mathsf{Impr}(\sqsubseteq_1, \sqsubseteq_2)$ are satisfied and yet, since a is in $\sqsubseteq_2$ at a lower level than in $\sqsubseteq_1$, then $P^{\top,\mathrm{eq}}(\sqsubseteq_1, \sqsubseteq_2)$ does not hold.

$\{Dom, SSCI\} \not\Longrightarrow \{SDomI\}$. Let P be the transitive closure of the union of $P^{\top,\perp}$ and P^{dom}. It is easy to verify that P is a preference satisfying $\{Dom, SSCI\}$. Consider the two rankings $\sqsubseteq_1$ and $\sqsubseteq_2$ such that $c \sqsubset_1 a$, $a \equiv_1 b$, $c \sqsubset_2 a$, and $c \equiv_2 b$. Since $\sqsubseteq_1$ and $\sqsubseteq_2$ are equivalent for $P^{\top,\perp}$, they are also equivalent for P. However, SDomI strictly prefers $\sqsubseteq_2$.

$\{DomI, SSw\} \not\Longrightarrow \{Dom\}$. Consider P^1 from Figure 2. It is straightforward to see that P^1 satisfies both DomI and SSw. To see that it violates Dom, consider the linear order $a \sqsubset_1 b \sqsubset_1 c$ and the ranking $a \sqsubset_2 b \equiv_2 c$. Since they are 3-tier equivalent for a, they are Dom-equivalent according to Lemma 1. However, P^1 strictly prefers $\sqsubseteq_2$.

$\{SCI, SSw\} \not\Longrightarrow \{SSCI, SSw\}$. Consider $P^{\top}$. It is immediate to see that $P^{\top}$ satisfies both SCI and SSw. Then, consider the rankings $\sqsubseteq_1$ and $\sqsubseteq_2$ over two agents such that $a \equiv_1 b$ and $b \sqsubset_2 a$. According to axiom SSCI $\sqsubseteq_2$ should be strongly preferable to $\sqsubseteq_1$, however they are equivalent for $P^{\top}$.

$\{Dom, SSw\} \not\Longrightarrow \{SSCI, Sw\}$. It is straightforward to see that $P^{\mathrm{ab}} = P^{\mathrm{bst}}$ satisfies both Dom and SSw. Consider the rankings $a \equiv_1 b$ and $b \sqsubset_2 a$, they are equivalent for P^{ab}, but SSCI requires $\sqsubseteq_2$ to be strongly preferred.

$\{SSCI, Sw\} \not\Longrightarrow \{SSw\}$. Consider the preference P^2 in Figure 3. It is straightforward to see that P^2 satisfies both SSCI and Sw. However, consider the rankings $\sqsubseteq_1$ and $\sqsubseteq_2$ such that $a \sqsubset_1 b$, $b \equiv_1 c$, $c \sqsubset_2 a$, and $a \equiv_2 b$. Axiom SSw would force $\sqsubseteq_2$ to be strongly preferable to $\sqsubseteq_1$, however they are equivalent for P^2. □

a,b,c	a,b c	a,c b	a b,c	a b c	a c b
		b a c	c a b		
	b,c a	b a,c	c a,b		
		b c a	c b a		

Figure 2. The preference relation P^1 used in Theorem 4. P^1 prefers the rankings that appear higher in the figure. This preference satisfies DomI, SSw, and Ind, but not Dom.

a b,c		
a,b,c	a b c	a c b
a,b c	a,c b	b,c a
b a c	c a b	
b a,c	c a,b	
b c a	c b a	

Figure 3. The preference relation P^2 used in Theorem 4. This preference satisfies SSCI, Sw, and Ind, but not SSw.

The next theorem characterizes the strongest preference theory $Pref_A(Dom, SDomI)$ in terms of P^{dom}. The proof is omitted due to space constraints.

Theorem 5 *A preference belongs to $Pref_A(Dom, SDomI)$ iff it refines and perfects P^{dom}.*

To appreciate the relevance of the previous theorem, notice that Lemma 1 tells us that all preferences satisfying the axiom Dom are 3-tier (that is, they do not distinguish between 3-tier equivalent rankings). However, if Dominance establishes a *strong* preference between two rankings, a preference satisfying Dom may instead equate those orders. For an extreme example, the degenerate preference $P^{\equiv}$ equates all rankings, but it still satisfies Dom. Theorem

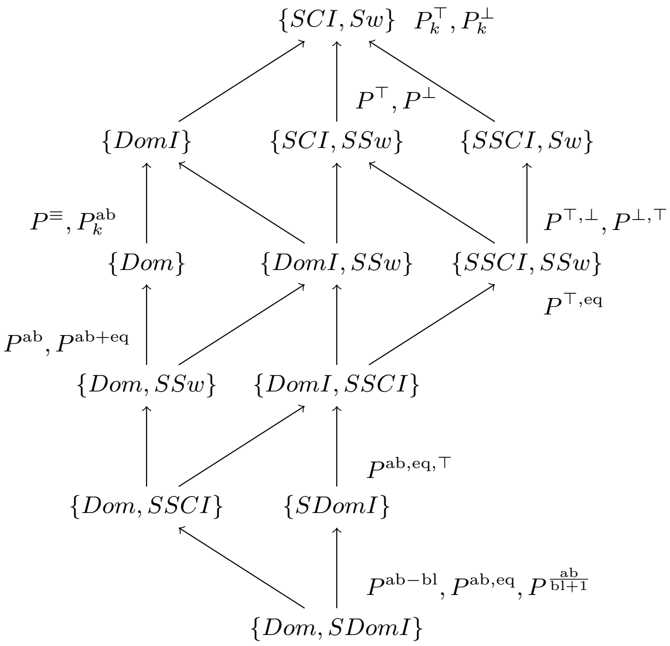

Figure 4. Hasse diagram for containment between self-interest theories. Lower theories are contained in higher ones. The preferences listed next to a theory belong to that theory and do not belong to any stronger theory in the diagram.

5 states that adding $SDomI$ to the picture forces preferences to uphold those strong preferences as well. Indeed, since preferences are total orders, a preference refining and perfecting P^{dom} can only (and must) establish a preference among rankings that are incomparable for P^{dom}.

Finally, the hierarchy depicted in Figure 4 does not change in case we assume indifference to the others.

Theorem 6 *Let $\mathcal{T}_1, \mathcal{T}_2 \subseteq \mathcal{AX}$. Then, $\mathcal{T}_1 \cup \{Ind\}$ implies $\mathcal{T}_2 \cup \{Ind\}$ if and only if $\mathcal{T}_1$ implies $\mathcal{T}_2$.*

Proof: The *if* direction being trivial, for the *only if* direction, assume by contraposition that $\mathcal{T}_1$ does not imply $\mathcal{T}_2$. By Theorem 3 we can assume that $\mathcal{T}_1$ and $\mathcal{T}_2$ are canonical. Now, all the non-implications in Theorem 4 have been proved using as witnesses preferences that satisfy Ind. Consequently, there exists a preference P that satisfies $\mathcal{T}_1 \cup \{Ind\}$ and does not satisfy $\mathcal{T}_2$. □

5 Quantitative Rewards

The designer of a competition may try to induce specific preferences by converting the outcome of the competition from a ranking to a numerical reward for each participant.

A *reward policy* (or simply *reward*) is a function $r : \mathcal{R}(A) \to \mathbb{R}^A$ preserving strict order in a weak sense; formally, for all $a, b \in A$, if $a \sqsubset b$ then $r(\sqsubseteq)(a) \leq r(\sqsubseteq)(b)$. We allow rewards to be more or less discriminating than the ranking. The only requirement is that they do not invert the order of the agents. A *faithful* reward preserves order: $a \sqsubseteq b$ iff $r(\sqsubseteq)(a) \leq r(\sqsubseteq)(b)$. Intuitively, a faithful reward only establishes the distance between different layers from the ranking, whereas a general reward can also perform tie-breaking. Additionally, a reward is *anonymous* if it is invariant under permutations of agents. Intuitively, anonymity is a fairness condition ensuring that the designer does not favour any agent in particular and the rewards are assigned only on the basis of the mutual placements in the ranking.

A straightforward application of the definition proves that anonymous rewards assign the same amount to equivalent agents. In other words, anonymous rewards cannot break ties.

Lemma 3 *If r is an anonymous reward and $\sqsubseteq$ is a ranking such that $a \equiv b$, then $r(\sqsubseteq)(a) = r(\sqsubseteq)(b)$.*

A reward r naturally induces a preference P_a^r for each agent $a \in A$, where $P_a^r(\sqsubseteq_1, \sqsubseteq_2)$ iff $r(\sqsubseteq_1)(a) \leq r(\sqsubseteq_2)(a)$. Moreover, we say that a preference is *compatible* with a reward r if it perfects the preference induced by r. We say that a reward r satisfies one of the preference axioms in Section 3 iff for all agents a, the preference P_a^r induced by r on a satisfies that axiom.

The next result draws a close connection between anonymity of a reward and the two axioms Sw and SSw.

Lemma 4 *If a reward is anonymous then it satisfies Sw. Moreover, an anonymous reward is faithful if and only if it satisfies SSw.*

A common requirement is that a reward should split a fixed amount of money (the jackpot) among the participants. In that case, we say that the reward is *cake-cutting* (CC) and we conventionally set the jackpot to 1. Formally, for all subsets of agents $A' \subseteq A$, denote by $r(\sqsubseteq)(A')$ the sum of the rewards $r(\sqsubseteq)(a)$ for all $a \in A'$. A reward is cake-cutting if for all rankings $\sqsubseteq$ it holds $r(\sqsubseteq)(A) = 1$.

5.1 Level-Averaged Rewards

Next, we introduce a class of anonymous and CC rewards, called *level-averaged*, that will be shown to play a special role in the taxonomy of rewards. Intuitively, level-averaged rewards generalize the so-called *fractional ranking* [15, 5], in that a group of tied agents receive the arithmetic mean of what they would receive if the tie was broken in an arbitrary way. The formal definition follows. Let $\sqsubseteq$ be a ranking with l levels, and n_i be the number of agents at level $i = 1, \dots, l$ (1 is the top level). Clearly, the sum n of all the n_i is equal to the total number of agents. A reward r is *level-averaged (LA)* iff there exists a sequence of real coefficients $\lambda_1 \geq \cdots \geq \lambda_n$ summing up to 1 such that if a is an agent at level i, it receives the reward

$$r(\sqsubseteq)(a) = \frac{1}{n_i} \sum_{j=N_i+1}^{N_i+n_i} \lambda_j,$$

where N_i is the sum of all n_h with $h < i$. A reward is *strictly level-averaged* if it is LA and its coefficients are strictly monotonic, that is, if $\lambda_1 > \cdots > \lambda_n$.

For example, consider a set of three agents a, b, c and the strictly LA reward r with coefficients $(\lambda_1, \lambda_2, \lambda_3) = (\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$. For any linear ranking, r attributes reward λ_i to the agent on level i. For the ranking $a \equiv b \sqsubset c$, instead, r assigns reward $\frac{1}{2}$ to c, because it is the only agent on level 1; agents a and b receive $\frac{1}{2}(\frac{1}{3} + \frac{1}{6}) = \frac{1}{4}$ each.

The first result on LA rewards shows which axioms they satisfy. In particular, strictly LA rewards induce preferences in the strongest of our theories.

Lemma 5 *Level-averaged rewards are anonymous, CC, and satisfy Dom. Strictly level-averaged rewards additionally satisfy $SDomI$.*

Proof: Anonymity and CC are obvious consequences of the definition of LA reward. As for Dom, let $\sqsubseteq_1$ and $\sqsubseteq_2$ be two rankings s.t. $\mathsf{Dom}_a(\sqsubseteq_1, \sqsubseteq_2)$ holds, for some agent a. Following the definition of LA reward, let n_1, m_1, n_2, m_2 be the integers such that:

$$r(\sqsubseteq_1)(a) = \frac{1}{n_1} \sum_{j=m_1}^{m_1+n_1-1} \lambda_j$$
$$r(\sqsubseteq_2)(a) = \frac{1}{n_2} \sum_{j=m_2}^{m_2+n_2-1} \lambda_j .$$

In words, $r(\sqsubseteq_1)(a)$ is the arithmetic mean of the sequence of values $\lambda_{m_1}, \ldots, \lambda_{m_1+n_1-1}$, and $r(\sqsubseteq_2)(a)$ is the arithmetic mean of $\lambda_{m_2}, \ldots, \lambda_{m_2+n_2-1}$. Recall that the sequence of λ_j's is monotonically non-increasing.

Now, inclusion (III) states that the set of agents above a in $\sqsubseteq_2$ is a subset of the set of agents above a in $\sqsubseteq_1$. This implies that $m_2 \leq m_1$. Additionally, inclusion (IV) states that the set of agents above or equivalent to a in $\sqsubseteq_2$ is also a subset of the corresponding set in $\sqsubseteq_1$. This implies that $m_2 + n_2 \leq m_1 + n_1$. Summarizing, the sequence of λ_j's whose mean is $r(\sqsubseteq_2)(a)$ starts earlier (index-wise) and ends earlier than the corresponding sequence for $r(\sqsubseteq_1)(a)$. It follows that $r(\sqsubseteq_2)(a) \geq r(\sqsubseteq_1)(a)$ and any a-preference compatible with r must weakly prefer $\sqsubseteq_2$ over $\sqsubseteq_1$.

Finally, we prove that if r is strictly LA, it additionally satisfies $SDomI$. Let $\sqsubseteq_1$ and $\sqsubseteq_2$ be two rankings s.t. $\mathsf{Dom}_a(\sqsubseteq_1, \sqsubseteq_2)$ and $\mathsf{Impr}_a(\sqsubseteq_1, \sqsubseteq_2)$ hold, for some agent a. From the previous argument, we know that $r(\sqsubseteq_2)(a) \geq r(\sqsubseteq_1)(a)$. It remains to show that Impr_a implies that such inequality is strict. By Lemma 1, one of the two inequalities $m_2 \leq m_1$ and $m_2 + n_2 \leq m_1 + n_1$ is strict. Thus, either the sequence of λ_j's whose mean is $r(\sqsubseteq_2)(a)$ starts strictly earlier than the corresponding sequence for $r(\sqsubseteq_1)(a)$, or the former ends strictly earlier than the latter. In either case, since the sequence of λ_j's is strictly decreasing, we obtain that $r(\sqsubseteq_2)(a) > r(\sqsubseteq_1)(a)$ and hence the thesis. □

We can actually prove that being a LA reward is equivalent to being anonymous, CC, and satisfying Dom.

Theorem 7 *A reward is level-averaged if and only if it is anonymous, CC, and satisfies Dom.*

Proof: The "only if" direction is a consequence of Lemma 5. The "if" direction is proved by induction on the number of agents n. The base case $n = 1$ trivially follows from the CC condition.

Assume by induction that any anonymous, CC, and Dom-satisfying reward on n agents is LA. Let r be a reward on $n+1$ agents $A \cup \{a\}$, which is anonymous, CC, and Dom-satisfying. We define r' as the projection of r on A obtained by fixing the position of a as the top element in the order. Formally, for all rankings $\sqsubseteq_A$ on A and agents $b \in A$, we set

$$r'(\sqsubseteq_A)(b) = \frac{r(\sqsubseteq_A^a)(b)}{r(\sqsubseteq_A^a)(A)},$$

where $\sqsubseteq_A^a$ is the ranking on $A \cup \{a\}$ obtained from $\sqsubseteq_A$ by placing a above all other agents.

The reward r' inherits anonymity from r and is CC by construction. We show that it also satisfies Dom. Let $\sqsubseteq_1$ and $\sqsubseteq_2$ be two rankings on A such that $\mathsf{Dom}_b(\sqsubseteq_1, \sqsubseteq_2)$ holds for an agent $b \in A$. Notice that adding a on top of two rankings preserves dominance among them from the point of view of all the other agents. Then, since r satisfies Dom, it holds (i) $r(\sqsubseteq_1^a)(b) \leq r(\sqsubseteq_2^a)(b)$. Moreover, since in both $\sqsubseteq_1^a$ and $\sqsubseteq_2^a$ the agent a is the only agent on level one, we have that $\mathsf{Dom}_a(\sqsubseteq_1^a, \sqsubseteq_2^a)$ and $\mathsf{Dom}_a(\sqsubseteq_2^a, \sqsubseteq_1^a)$, which implies (ii) $r(\sqsubseteq_1^a)(a) = r(\sqsubseteq_2^a)(a)$. In turn, conditions (i) and (ii) imply that

$$\begin{aligned} r'(\sqsubseteq_1)(b) &= \frac{r(\sqsubseteq_1^a)(b)}{r(\sqsubseteq_1^a)(A)} && \text{by definition} \\ &\leq \frac{r(\sqsubseteq_2^a)(b)}{r(\sqsubseteq_1^a)(A)} && \text{by (i)} \\ &= \frac{r(\sqsubseteq_2^a)(b)}{r(\sqsubseteq_2^a)(A)} && \text{by (ii)} \\ &= r'(\sqsubseteq_2)(b) && \text{by definition,} \end{aligned}$$

and hence r' satisfies Dom. Summarizing, we have that r' is anonymous, CC, and satisfies Dom. Therefore, by the induction hypothesis r' is LA with coefficients, say, $\{\lambda_1, \ldots, \lambda_n\}$.

Let $\sqsubseteq_A$ be an arbitrary ranking over A and $\sqsubseteq_A^a$ be the extension where a is placed above all other agents. We set λ_0 to $r(\sqsubseteq_A^a)(a)$ and λ_i' to $\lambda_i \cdot r(\sqsubseteq_A^a)(A)$. Notice that by condition (ii) above λ_0 does not depend on the choice of $\sqsubseteq_A$. Clearly, it holds by construction that $\lambda_1' \geq \ldots \geq \lambda_n'$ and $\lambda_0 + \lambda_1' + \cdots + \lambda_n' = 1$. Moreover, we show that $\lambda_0 \geq \lambda_1'$. Let $\sqsubseteq$ be a linear order on $A \cup \{a\}$ with a the only agent on level 1. Since we have more than one agent, there exists an agent b on level 2 in $\sqsubseteq$. Consider the linear order $\sqsubseteq'$ obtained from $\sqsubseteq$ by swapping a and b. We have that

$$\begin{aligned} r(\sqsubseteq)(a) &\geq r(\sqsubseteq')(a) && \text{by } Dom \\ &= r(\sqsubseteq)(b) && \text{by anonymity.} \end{aligned}$$

Since by construction $\sqsubseteq$ is equal to $\sqsubseteq_A^a$ for some linear order $\sqsubseteq_A$, we know that $r(\sqsubseteq)(a) = \lambda_0$ and $r(\sqsubseteq)(b) = \lambda_1'$. Therefore, $\lambda_0 \geq \lambda_1'$ and $\lambda_0, \lambda_1', \ldots, \lambda_n'$ have all the properties required to form a LA reward. Indeed, we now prove that r is LA with coefficients $\lambda_0, \lambda_1', \ldots, \lambda_n'$.

Let $\sqsubseteq$ be a ranking on $A \cup \{a\}$, we can assume w.l.o.g. that agent a is at level one — otherwise we can consider by anonymity an r-equivalent ranking obtained by swapping a with some agent on the top. Let $\sqsubseteq^+$ be obtained from $\sqsubseteq$ by moving a above all other agents.

First, consider the case where a is the only agent on the top level in $\sqsubseteq$. Then $\sqsubseteq^+ = \sqsubseteq$ and $\sqsubseteq$ is equal to $\sqsubseteq_A^a$ for some ranking $\sqsubseteq_A$ on A. As shown previously, this means that $r(\sqsubseteq)(a) = \lambda_0$. Moreover, given an agent $b \in A$ located on level $i > 1$, it holds

$$\begin{aligned} r(\sqsubseteq)(b) &= r(\sqsubseteq_A^a)(b) = r'(\sqsubseteq_A)(b) \cdot r(\sqsubseteq_A^a)(A) \\ &= \frac{r(\sqsubseteq_A^a)(A)}{n_i} \sum_{j=N_i}^{N_i+n_i-1} \lambda_j = \frac{1}{n_i} \sum_{j=N_i}^{N_i+n_i-1} \lambda_j' . \end{aligned}$$

Consequently, r is LA on $\sqsubseteq$.

Assume now that a is *not* the only agent on the top level in $\sqsubseteq$ (i.e., $\sqsubseteq^+ \neq \sqsubseteq$). By the previous case we know that r is LA on $\sqsubseteq^+$. In particular, $r(\sqsubseteq^+)(a) = \lambda_0$ and, assuming that the number of agents on the second level of $\sqsubseteq^+$ is k, for each of those agents b it holds $r(\sqsubseteq^+)(b) = \frac{1}{k}(\lambda_1' + \cdots + \lambda_k')$. Let c be an agent at level ≥ 2 in $\sqsubseteq$ (if any). By Lemma 1, $\sqsubseteq$ and $\sqsubseteq^+$ are Dom-equivalent for c and hence $r(\sqsubseteq)(c) = r(\sqsubseteq^+)(c)$. Since this holds for all agents below level one in $\sqsubseteq$, then the total budget that r can assign to the first level in $\sqsubseteq$ is equal to the budget for the first two levels in $\sqsubseteq^+$ which is $\lambda_0 + \lambda_1' + \cdots + \lambda_k'$. Finally, by anonymity, when evaluated on $\sqsubseteq$ r must equally distribute that same budget over agent a and the other k agents. Consequently, each of them receives a reward equal to $\frac{1}{k+1} \cdot (\lambda_0 + \lambda_1' + \cdots + \lambda_k')$, which means that r is LA. □

Similarly to Theorem 7, we can prove that for a reward to be anonymous, CC and satisfying Dom and SSw is equivalent to being strictly LA. Since strictly LA rewards additionally satisfy $SDomI$ (Lemma 5), this implies that the theories $\{Dom, SSw\}$, $\{Dom, SSCI\}$, and $\{Dom, SDomI\}$ collapse when applied to rewards. In other words, there is no reward that induces a preference satisfying $\{Dom, SSw\}$ but not $SDomI$, because all preferences that satisfy the former combination of axioms are strictly LA.

Theorem 8 *A reward is strictly level-averaged if and only if it is anonymous, CC, and satisfies* Dom *and* SSw*.*

The taxonomy of anonymous and CC rewards w.r.t. the axioms in $\mathcal{AX}$ is depicted in Figure 5. The implications are all consequences of Theorem 4 (see Figure 4). The non-implications need new counter-examples, because the preferences used to prove non-implications in Theorem 4 are not necessarily induced by a reward. To find these new counter-examples, we encoded the existence of a suitable reward as a system of linear inequalities and used the Parma Polyhedra Library [2] to check its feasibility. The details are omitted due to space limitations.

Theorem 9 *Let* $\mathcal{T} \subseteq \mathcal{AX}$ *be a set of axioms and let* $Rew(\mathcal{T})$ *be the set of anonymous and CC rewards satisfying the axioms in* $\mathcal{T}$*. If* $Rew(\mathcal{T}) \subseteq Rew(\{SCI, Sw\})$*, then* $Rew(\mathcal{T})$ *coincides with one of the classes* $Rew(\mathcal{T}')$*, with* $\mathcal{T}'$ *in Figure 5. Additionally, a class* $Rew(\mathcal{T}_1)$ *is contained in another class* $Rew(\mathcal{T}_2)$ *if and only if there is a path from* $\mathcal{T}_1$ *to* $\mathcal{T}_2$ *in that figure.*

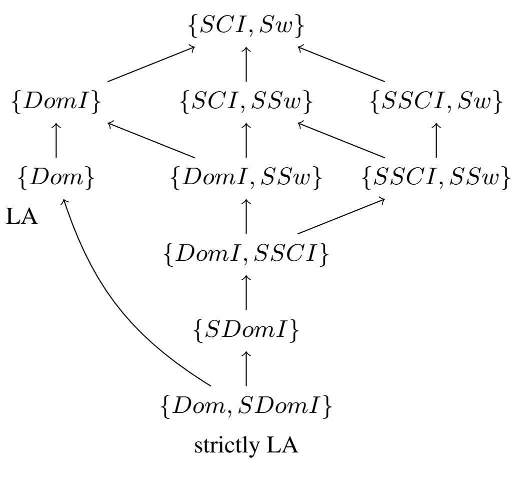

Figure 5. Hasse diagram for containment between classes of anonymous and CC rewards.

6 Conclusions

This paper initiates an investigation about the rational behavior of self-interested agents participating in a competition whose outcome is a ranking (i.e., a weak order) among participants. Even though several works focused on preferences over aggregate outcomes (see [3] and references therein), studying preferences over rankings is to the best of our knowledge a novel research question.

Our results show an intriguing landscape of self-interest theories, most of which include simple preferences that are easily motivated by practical scenarios. From the point of view of the competition designer, being able to pinpoint the likely preference theory adopted by the participants is valuable information because it allows the designer to predict the participants' behavior during the competition. For example, in [7] the authors address the problem of avoiding irrelevant matches in sporting tournaments by assuming that agents adopt the preference $P^{\mathrm{ab,eq}}$.

Generally speaking, the designer will want participants to act competitively, rather than indifferently or even cooperatively. In this respect, we investigated the case when the competition designer is able to assign monetary rewards to the participants, based on the weak order. Even if the assigned rewards do not break the ties in the order, they may still influence the agents' preferences simply by tuning the *distance* between different levels in the weak order. In particular, we have characterized a class of level-averaged rewards that are anonymous, cake-cutting, and promote competition by inducing agents to overcome each other.

REFERENCES

[1] Ernest W Adams, 'Elements of a theory of inexact measurement', *Philosophy of science*, **32**(3/4), 205–228, (1965).
[2] R. Bagnara, P. M. Hill, and E. Zaffanella, 'The Parma Polyhedra Library: Toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems', *Science of Computer Programming*, **72**(1–2), 3–21, (2008).
[3] Salvador Barberà, Walter Bossert, and Prasanta K. Pattanaik, 'Ranking sets of objects', in *Handbook of Utility Theory: Volume 2 Extensions*, eds., Salvador Barberà, Peter J. Hammond, and Christian Seidl, chapter 17, 893–977, Springer US, Boston, MA, (2004).
[4] Markus Brill, Piotr Faliszewski, Frank Sommer, and Nimrod Talmon, 'Approximation algorithms for balancedCC multiwinner rules', in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS19), Montreal, QC, Canada, May 13-17, 2019*, pp. 494–502, (2019).
[5] Hanna Bury and Dariusz Wagner, 'Group judgement with ties. Distance-based methods', in *New Approaches in Automation and Robotics*, chapter 9, 153–172, IntechOpen, (05 2008).
[6] E. Elkind and P. Faliszewski, 'Properties of multiwinner voting rules', *Social Choice and Welfare*, **48**(3), 599–632, (2017).
[7] M. Faella and L. Sauro, 'Do all tournaments admit irrelevant matches?', in *Proc. of the 17th Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS18)*, pp. 982–989, (2018).
[8] Luis Sánchez Fernández and Jesús A. Fisteus, 'Monotonicity axioms in approval-based multi-winner voting rules', in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS19), Montreal, QC, Canada, May 13-17, 2019*, pp. 485–493, (2019).
[9] R. Freeman, M. Brill, and V. Conitzer, 'General tiebreaking schemes for computational social choice', in *Proc. of the 17th Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS15), Istanbul, Turkey, May 4-8*, pp. 1401–1409, (2015).
[10] J. Lang, J. Monnot, and A. Slinko, 'Beyond electing and ranking: Collective dominating chains, dominating subsets and dichotomies', in *Proc. of the International conference on Autonomous Agents and Multi-Agent Systems (AAMAS17), Sao Paulo, Brazil, May 812*, (2017).
[11] R. Myerson, 'Axiomatic derivation of scoring rules without the ordering assumption', *Social Choice and Welfare*, **21**(12), 59–74, (1995).
[12] Tao Qin, Wei Chen, and Tie-Yan Liu, 'Sponsored search auctions: Recent advances and future directions', *ACM Trans. Intell. Syst. Technol.*, **5**(4), 60:1–60:34, (January 2015).
[13] M. Remzi Sanver and William S. Zwicker, 'Monotonicity properties and their adaptation to irresolute social choice rules', *Social Choice and Welfare*, **39**(2), 371–398, (Jul 2012).
[14] L. Sauro, 'On the hierarchical nature of partial preferences over lotteries', *Autonomous Agents and Multi-Agent Systems*, **31**(6), 1467–1505, (2017).
[15] Lawrence Seiford and Wade Cook, 'Priority ranking and consensus formation', *Management Science*, **24**, 1721–1732, (12 1978).