

Fair Allocation Based on Diminishing Differences

Erel Segal-Halevi
 Bar-Ilan University
 erelsgl@gmail.com

Haris Aziz
 Data61 CSIRO and UNSW
 haris.aziz@data61.csiro.au

Avinatan Hassidim
 Bar-Ilan University
 avinatan@cs.biu.ac.il

Abstract

Ranking alternatives is a natural way for humans to explain their preferences. It is being used in many settings, such as school choice (NY, Boston), course allocations, and the Israeli medical lottery. In some cases (such as the latter two), several “items” are given to each participant. Without having any information on the underlying cardinal utilities, arguing about fairness of allocation requires extending the ordinal item ranking to ordinal bundle ranking. The most commonly used such extension is stochastic dominance (SD), where a bundle X is preferred over a bundle Y if its score is better according to all additive score functions. SD is a very conservative extension, by which few allocations are necessarily fair while many allocations are possibly fair.

We propose to make a natural assumption on the underlying cardinal utilities of the players, namely that the difference between two items at the top is at least as large as the difference between two items down the list. This assumption implies a preference extension which we call diminishing differences (DD), where a X is preferred over Y if its score is better according to all additive score functions satisfying the DD assumption.

We give a full characterization of allocations that are necessarily-proportional or possibly-proportional according to this assumption. Based on this characterization, we present a polynomial-time algorithm for finding a necessarily-DD-proportional allocation if it exists. Simulations based on a simple random model show that with high probability, a necessarily-proportional allocation does not exist but a necessarily-DD-proportional allocation exists. Moreover, that allocation is proportional according to the underlying cardinal utilities.

1 Introduction

Algorithms for fair assignment of indivisible items differ in the kind of information they require from the users.

Some algorithms require the users to rank bundles of items (i.e, report a total order among the bundles). Examples are the Decreasing Demand procedure [Herreiner and Puppe, 2002], Approximate-CEEI procedure [Budish, 2011] and Undercut two-agent procedure [Brams *et al.*, 2012; Aziz, 2015]. The computational and communicational burden might be large, since the number of bundles is exponential in the number of items.

Other algorithms require the users to evaluate individual items (i.e, supply a numeric monetary value for each item). Such algorithms are often termed *cardinal*. They often assume that the users’ valuations are additive, so that the value of a bundle can be calculated by summing the values of the individual items. Examples are the Adjusted Winner procedure [Brams and Taylor, 2000], approximate-maximin-share procedure [Procaccia and Wang, 2014] and Maximum Nash Welfare procedure [Caragiannis *et al.*, 2016]. In this setting the communication is linear in the number of items, but the mental burden may still be large, since assigning an exact monetary value to individual items is not easy. This is especially true when items are valued for personal reasons (such as when dividing inheritance) and do not have a market price.

This paper focuses on a third class of algorithms, which only require the users to rank (i.e, report a total order among) individual items. Such algorithms are often termed *ordinal*.

Ordinal algorithms are ubiquitous in mechanism design, and are often used in real world applications, such as the National Residency Matching Program [Roth and Peranson, 1997] (even when married couples insert their preferences together [Ashlagi *et al.*, 2014]), school choice applications [Abdulkadiroglu and Sönmez, 2003], and university admittance [Hassidim *et al.*, 2016a; 2016b]. One reason for this is that it is relatively easy for people to state ordinal preferences. Another reason is related to legacy systems: often the designer can change the allocation mechanism, but can not change the input procedure, as agents do not want to learn new ways to insert their input to the system.

Ordinal algorithms are also common in AI and in fair division. Examples are the AL two-agent procedure [Brams *et al.*, 2013], optimal-proportional procedure [Aziz *et al.*, 2015], picking-sequence procedures [Brams and Kaplan, 2004; Bouveret and Lang, 2011] and the envy-free procedures of [Bouveret *et al.*, 2010]. Such algorithms often assume that the agents’ preferences are implicitly represented by an additive

utility function, which is not known to the algorithm. This creates ambiguity in the agents’ bundle rankings. For example, if an agent ranks four items as $w \succ x \succ y \succ z$, then, based on additivity, the algorithm can know that e.g. $\{w, x\} \succ \{y, z\}$ and $\{w, y\} \succ \{x, z\}$, but cannot know the relation between $\{w, z\}$ and $\{x, y\}$. Algorithms cope with this problem in several ways.

1. **Necessary-fairness criteria.** An allocation is called **necessarily-fair** if it is fair for *all* additive utility profiles consistent with the reported item-rankings. Here, “fair” may be substituted by any fairness criterion, such as envy-freeness or proportionality, as well as Pareto-efficiency. Necessary fairness is a strong requirement, which is not always satisfiable. For example, the AL procedure finds a necessarily-envy-free allocation, but only for two agents, and even then, it might need to discard some of the items in order to keep its guarantee for necessary-envy-freeness.

2. **Possible-fairness criteria.** An allocation is called **possibly-fair** if it is fair for *at least one* additive utility profile consistent with the reported item-rankings. Again, “fair” may be substituted by proportional or envy-free or Pareto-efficient. Possible fairness is a weak criterion; algorithms that only return possibly-fair allocations might be considered unfair by users whose actual utility function is different than the one assumed by the algorithm.

3. **Scoring rules.** A scoring rule is a function that maps the rank of an item to a numeric score. A common example is the Borda scoring rule [Young, 1974], where the least desired item has a score of 1, the next item has a score of 2, and so on. The score of a bundle is the sum of the scores of its items. It is assumed that all agents have the same scoring function. I.e., even though agents may rank items differently, the mapping from the ranking to the numeric utility function is the same for all agents [Bouveret and Lang, 2011; Kalinowski *et al.*, 2013; Baumeister *et al.*, 2016; Darmann and Klamlner, 2016]. This strong assumption weakens the fairness guarantee. Allocation may appear unfair to agents whose actual scoring rule is different.

1.1 Contribution

The present paper suggests an alternative between the strong guarantee of necessary-fairness and the weak guarantee of possible-fairness and scoring-rule-fairness.

We assume that people care more about their high-valued items than about their lower-valued items. Specifically, we assume that the utility-difference between the best item and the second-best item is at least as large as the utility between the second-best and the third-best, and so on. We call this assumption *Diminishing Differences (DD)*. The DD assumption is satisfied by the Borda scoring rule, as well as by many other scoring rules, as well as by lexicographic preferences.

DD is justified in many settings where the agents are more concerned about getting a most preferred item than about not getting a least preferred item. For example, in a matching of doctors to internships, it was reported that doctors care the most about being assigned to one of their top choices [Bronfman *et al.*, 2015a; 2015b].

Based on the DD assumption, we formalize several fairness notions. We call an allocation **necessarily-DD-fair** (NDD-

fair) if it is fair according to all additive utility profiles satisfying the DD assumption, and **possibly-DD-fair** (PDD-fair) if it is fair according to at least one additive utility profile satisfying the DD assumption. Again, “fair” may be substituted by envy-free or proportional or Pareto-efficient. The following implications are obvious for any fairness criterion:

$$\text{Nec-fair} \implies \text{NDD-fair} \implies \text{PDD-fair} \implies \text{Pos-fair}$$

In other words, the DD-fairness criteria are intermediate in strength between necessary-fairness and possible-fairness. A formal definition of these concepts appears in Section 3.

The first question of interest is to decide, given an item ranking and two bundles, whether the NDD or the PDD relation holds between these bundles. We prove characterizations of the NDD and PDD set relations that provide linear-time algorithms for answering these questions (Section 4).

Next, we prove a necessary and sufficient condition for the existence of an NDD-proportional (NDDPR) allocation. Essentially, an NDDPR allocation exists iff it is possible to (a) give all agents the same number of items and (b) give each agent his best item. The proof is constructive and presents a simple linear-time algorithm for finding an NDDPR allocation if it exists (Section 5).

It is interesting to compare the above condition to Condition D of Brams *et al.* [2013], which is necessary and sufficient for the existence of a NecPR allocation for two agents. For NecPR, it is required that for every odd integer $k \in \{1, 3, \dots, 2n - 1\}$, the agents have a different set of k best items; in particular, they should have a different worst item (see Example 2 in Sec. 3). For NDDPR, it is only required that the agents have a different best item. This means that an algorithm that returns NDDPR allocations may have larger ‘recall’ than an algorithm that returns only NecPR allocations. On the other hand, the ‘precision’ of such an algorithm might be lower (i.e., its output allocations may be considered unfair by some agents). To assess the magnitude of these effects, we present a simple preliminary simulation experiment. We construct partially-correlated utility profiles at random, estimate the probability that an NDDPR/NecPR allocation exists, and check whether the NDDPR allocation given by our linear-time algorithm is indeed proportional according to the underlying cardinal utilities. We find that the increase in recall is substantial and ranges between 20% and 40%, but the decrease in precision is not substantial: when there are sufficiently many items, our simple NDDPR algorithm almost always results in an allocation which is proportional according to the cardinal utilities. While this experiment is preliminary, it indicates that there is potential for further investigation of NDDPR as a normative fairness criterion (Section 6).

While our focus on this paper is on proportionality, we have results on other criteria, namely envy-freeness and Pareto-efficiency (Section 7). Whereas the DD assumption has a substantial effect on fairness notions, we show that the DD assumption does not lead to a new notion of Pareto-efficiency (PE). We show that NDD-PE is equivalent to necessary-PE and PDD-PE is equivalent to possible-PE. Finally, we show that checking whether an NDDEF allocation exists is NP-complete.

1.2 Related Work

Extending preferences over individual objects to sets of objects is a natural and principled way of succinctly encoding preferences [Barberà *et al.*, 2004]. One of the most common set-extensions is *stochastic dominance* (SD). It was developed for a different but related problem — extending preferences over individual outcomes to *lotteries* on outcomes. If X, Y are lotteries, then $X \succeq^{SD} Y$ iff $E[u(X)] \geq E[u(Y)]$ for every weakly-increasing utility function u [Hadar and Russell, 1969; Brandt, 2017]. In the context of fair item allocation, SD leads to the notions of necessary-fairness and possible-fairness [Aziz *et al.*, 2015]. Other common extensions are *downward-lexicographic* (DL) and *upward-lexicographic* (UL) [Cho, 2012; Bouveret *et al.*, 2010; Nguyen *et al.*, 2015].

The DD extension, which is the focus of this paper, is quite natural but has not been formalized in prior work. Interestingly, the DD extension is closely related to *second-order stochastic dominance* (SSD). If X, Y are lotteries, then $X \succeq^{SSD} Y$ iff $E[u(X)] \geq E[u(Y)]$ for every utility function u which is weakly-increasing and *weakly-concave* [Hadar and Russell, 1969]. In the context of item assignment, weak concavity is equivalent to *increasing differences* — agents care more about *not* getting the *worst* item than about getting the best item. Increasing differences make sense in fair division of chores [Aziz *et al.*, 2017]; in future work we plan to study which of our results apply to that setting.

Besides fair division, set-extensions have been applied for committee voting [Aziz *et al.*, 2016a] and social choice correspondences (see e.g., [Barberà *et al.*, 2001]).

In social choice theory, it is common to study restricted domains of preference profiles, such as single-peaked, single-crossing or level- r -consensus [Mahajne *et al.*, 2015; Nitzan *et al.*, 2017]. Many problems are much easier to solve in such restricted domains than in the domain of all preferences [Elkind and Lackner, 2014; Elkind *et al.*, 2017]. The present paper focuses on a restriction to preferences satisfying the DD assumption, which has not been studied so far.

2 Preliminaries

There is a set \mathcal{N} of agents with $n = |\mathcal{N}|$. There is a set \mathcal{M} of distinct items with $M = |\mathcal{M}|$. A *bundle* is a set of items. A *multi-bundle* is a multi-set of items, i.e, it may contain several copies of the same item.

An *allocation* X is a function that assigns to each agent i a bundle X_i , such that $\mathcal{M} = X_1 \cup \dots \cup X_n$ and the X_i -s are pairwise-disjoint. Each agent i has a strict ranking \succ_i on items. Each agent may also have a utility function u_i on (multi-)bundles. All utility functions considered in this paper are positive and additive, so the utility of a (multi-)bundle is the sum of the utilities of the items in it. A utility function u_i is *consistent with* \succ_i on items if for every two items x, y : $u_i(\{x\}) > u_i(\{y\}) \iff x \succ_i y$. The set of all additive utility functions consistent with \succ_i is denoted by $\mathcal{U}(\succ_i)$. The following definition is well-known:

Definition 1. For (multi-)bundles X_i, Y_i :

$$\begin{aligned} X_i \underset{\sim_i}{\sim_i}^{Nec} Y_i &\iff \forall u_i \in \mathcal{U}(\succ_i): u_i(X_i) \geq u_i(Y_i) \\ X_i \underset{\sim_i}{\sim_i}^{Pos} Y_i &\iff \exists u_i \in \mathcal{U}(\succ_i): u_i(X_i) \geq u_i(Y_i) \end{aligned}$$

3 DD Utility Functions and Relations

Given a strict ranking \succ_i , we assign to each item x a *level*, denoted $\text{LEV}_i(x)$, such that the level of the best item is M , the level of the second-best item is $M - 1$, etc (this is also known as the *Borda score* of the item).

Definition 2. Let \succ_i be a preference relation and u_i a utility function consistent with \succ_i . u_i has the *Diminishing Differences* (DD) property if, for every three items with consecutive levels $x \succ_i y \succ_i z$ such that $\text{LEV}_i(x) = \text{LEV}_i(y) + 1 = \text{LEV}_i(z) + 2$, it holds that $u_i(x) - u_i(y) \geq u_i(y) - u_i(z)$.

The set of all additive DD utility functions consistent with \succ_i is denoted by $\mathcal{U}^{DD}(\succ_i)$. The Borda utility function is a member of \mathcal{U}^{DD} . Another example of a member in \mathcal{U}^{DD} is the function $\text{Lex}_i(x) := 2^{\text{LEV}_i(x)}$, by which bundles are ordered by whether they contain the best item, then by whether they contain the second-best item, etc. An equivalent definition of \mathcal{U}^{DD} is given by the following lemma. The proof is arithmetic and it is omitted.

Lemma 1. The following are equivalent:

- (i) $u_i \in \mathcal{U}^{DD}(\succ_i)$
- (ii) For every four items x_1, x_2, x_3, x_4 with $x_1 \succeq_i x_3$ and $x_2 \succeq_i x_4$ and $x_1 \neq x_2$ and $x_3 \neq x_4$:

$$\frac{u_i(x_1) - u_i(x_2)}{\text{LEV}_i(x_1) - \text{LEV}_i(x_2)} \geq \frac{u_i(x_3) - u_i(x_4)}{\text{LEV}_i(x_3) - \text{LEV}_i(x_4)}$$

3.1 Diminishing-Differences Preference Relations

We are now ready to present our main objects of interest — the Diminishing-Differences relations.

Definition 3. For (multi-)bundles X_i, Y_i :

$$\begin{aligned} X_i \underset{\sim_i}{\sim_i}^{NDD} Y_i &\iff \forall u_i \in \mathcal{U}^{DD}(\succ_i): u_i(X_i) \geq u_i(Y_i) \\ X_i \underset{\sim_i}{\sim_i}^{PDD} Y_i &\iff \exists u_i \in \mathcal{U}^{DD}(\succ_i): u_i(X_i) \geq u_i(Y_i) \end{aligned}$$

Remark 1. Comparing Definitions 1 and 3, it is clear that:

$$X_i \underset{\sim_i}{\sim_i}^{Nec} Y_i \implies X_i \underset{\sim_i}{\sim_i}^{NDD} Y_i \implies X_i \underset{\sim_i}{\sim_i}^{PDD} Y_i \implies X_i \underset{\sim_i}{\sim_i}^{Pos} Y_i$$

3.2 Diminishing-Differences Fairness

For every integer k and bundle X_i , define $k \cdot X_i$ as the multi-bundle in which each item of X_i is copied k times. We define proportionality by comparing, for each agent i , the bundle X_i copied n times, to the bundle of all items \mathcal{M} .

Definition 4 (DD-proportionality). An allocation X is called

- Necessary-DD-proportional (NDDPR) if: $\forall i \in \mathcal{N} : n \cdot X_i \underset{\sim_i}{\sim_i}^{NDD} \mathcal{M}$.
- Possible-DD-proportional (PDDPR) if: $\forall i \in \mathcal{N} : n \cdot X_i \underset{\sim_i}{\sim_i}^{PDD} \mathcal{M}$.

Definition 5 (DD-envy-freeness). An allocation X is called

- Necessary-DD-envy-free (NDDEF) if for all agents i, j : $X_i \underset{\sim_i}{\sim_i}^{NDD} X_j$.

- Possible-DD-envy-free (PDDEF) if for all agents i, j : $X_i \succsim_i^{PDD} X_j$.

We now show that the DD assumption significantly affects the necessary and possible fairness requirements and results in meaningful fairness concepts. We give examples with two agents, so proportionality and envy-freeness coincide.

Example 1. Alice and Bob have the same preferences: $2m \succ 2m - 1 \succ \dots \succ 4 \succ 3 \succ 2 \succ 1$, for some $m \geq 3$. Both Alice and Bob get m items: Alice gets $2m, 2m - 1, \dots, m + 3, m + 2, 1$ and Bob gets $m + 1, m, \dots, 3, 2$. Intuitively this allocation seems very unfair, since Alice gets all the $m - 1$ best items. However, it is possibly-fair, since Bob's utility function might assign the a value near 0 to item 1 and a value near 1 to all other items. In better accordance with our intuition, the above allocation is not PDD-fair: by Theorem 2 below, Bob's bundle is not PDD-better than Alice's bundle, since Alice's bundle is NDD-better and its level is strictly higher.

Example 2. Now, Alice and Bob have almost opposite preferences, for some $m \geq 2$:

Alice: $2m \succ 2m - 1 \succ \dots \succ 4 \succ 3 \succ 2 \succ 1$.

Bob: $2 \succ 3 \succ 4 \succ \dots \succ 2m - 1 \succ 2m \succ 1$.

Intuitively we would expect that opposite preferences make it easy to attain a fair division. However, in this case no necessarily-fair allocation exists. This is because Alice and Bob have the same worst item (1), and under the necessary-fairness assumption, no one agrees to receives this item (since the utility function of the agent who receives this item might assign a value near 0 to this item and a value near 1 to all other items). In contrast, our Theorem 3 shows that an NDD-fair allocation exists. Intuitively, since it is possible to give each agent his/her best items, they are willing to compromise on the less important items.

4 Characterizing NDD and PDD Relations

In this section we are given a preference relation \succ_i on items and two multi-bundles X, Y , and have to decide whether $X \succsim_i^{NDD} Y$. We begin by proving a convenient characterization of the NDD relation. For the characterization, we define the *level* of a multi-bundle as the sum of the levels of the items in it: $\text{LEV}(X) := \sum_{x \in X} \text{LEV}(x)$ (this is equal to the Borda score of X . Note that all copies of the same item have the same level). We order the items in each multi-bundle by decreasing level, so $X = \{x_1, \dots, x_{|X|}\}$ where $x_1 \succeq_i \dots \succeq_i x_{|X|}$ (the order between different copies of the same item is arbitrary). For each $k \leq |X|$ we define X_k as the k best items in X , $X_k := \{x_1, \dots, x_k\}$.

Theorem 1. Let X, Y be two multi-bundles. $X \succsim_i^{NDD} Y$ iff

- (i) $|X| \geq |Y|$ and
- (ii) for each $k \in \{1, \dots, |Y|\}$: $\text{LEV}_i(X_k) \geq \text{LEV}_i(Y_k)$.

Proof. \implies : We have to prove that if (i) $|Y| > |X|$, or (ii) $\text{LEV}(Y_k) > \text{LEV}(X_k)$ for some k , then $X \not\sucsim_i^{NDD} Y$, i.e., there is a DD utility function u_i such that $u_i(X) < u_i(Y)$.

- (i) If $|Y| > |X|$, then take $u_i(x) = M^2 + \text{LEV}_i(x)$. The term M^2 is so large that the utility of a bundle is dominated by its cardinality, so $u_i(X) < u_i(Y)$.

- (ii) If for some k , $\text{LEV}_i(Y_k) > \text{LEV}_i(X_k)$, then let k be the smallest integer that satisfies this condition; hence $y_k \succ_i x_k$. Let $C := \text{LEV}_i(x_k) - 1$ and define u_i as:

$$\begin{aligned} u_i(x) &= [\text{LEV}_i(x) - C] \cdot M^2 & \text{for } x \succeq_i x_k \\ u_i(x) &= \text{LEV}_i(x) & \text{for } x_k \succ_i x \end{aligned}$$

The term M^2 is so large that the utility of a bundle is dominated by the level of its items that are weakly better than x_k , so again $u_i(X) < u_i(Y)$.

\Leftarrow : Assume that $X \not\sucsim_i^{NDD} Y$, and let u_i be an additive function in $\mathcal{U}^{DD}(\succ_i)$ for which $u_i(Y) > u_i(X)$. We have to prove that either (i) $|Y| > |X|$ or (ii) $\text{LEV}_i(Y_k) > \text{LEV}_i(X_k)$ for some k . There are two cases.

Case #1: $|Y| > |X|$. Then (i) is satisfied and we are done.

Case #2: $|Y| \leq |X|$. Then, since $u_i(Y) > u_i(X)$, there must be some $k \in \{1, \dots, |Y|\}$ for which $u_i(Y_k) > u_i(X_k)$.

Consider the *smallest* k for which this inequality holds for some $u_i \in \mathcal{U}^{DD}(\succ_i)$. So for $j < k$, there exists no function in $\mathcal{U}^{DD}(\succ_i)$ which assigns to Y_j a higher value than to X_j . This means that $u_i(y_k) > u_i(x_k)$, and the difference is substantial enough to make the value of Y_k more than the value of X_k , i.e:

$$u_i(y_k) - u_i(x_k) > \sum_{j=1}^{k-1} (u_i(x_j) - u_i(y_j)) \quad (1)$$

By our ordering of the items, $\forall j < k$ it holds that $y_j \succeq_i y_k$ and $x_j \succeq_i x_k$. Hence, by Lemma 1, whenever $x_j \neq y_j$:

$$\begin{aligned} \frac{\text{LEV}_i(y_k) - \text{LEV}_i(x_k)}{u_i(y_k) - u_i(x_k)} &\geq \frac{\text{LEV}_i(y_j) - \text{LEV}_i(x_j)}{u_i(y_j) - u_i(x_j)} \\ &= \frac{\text{LEV}_i(x_j) - \text{LEV}_i(y_j)}{u_i(x_j) - u_i(y_j)} \end{aligned} \quad (2)$$

so:

$$\frac{\text{LEV}_i(y_k) - \text{LEV}_i(x_k)}{u_i(y_k) - u_i(x_k)} \geq \frac{\sum_{j=1}^{k-1} (\text{LEV}_i(x_j) - \text{LEV}_i(y_j))}{\sum_{j=1}^{k-1} (u_i(x_j) - u_i(y_j))} \quad (3)$$

(the right-hand side of (3) is a weighted average of the right-hand sides of (2)).

Combining (1) and (3) gives:

$$\begin{aligned} \text{LEV}_i(y_k) - \text{LEV}_i(x_k) &> \sum_{j=1}^{k-1} (\text{LEV}_i(x_j) - \text{LEV}_i(y_j)) \\ \implies \text{LEV}_i(Y_k) &> \text{LEV}_i(X_k), \end{aligned}$$

so condition (ii) is satisfied and we are done. \square

Example 3. Consider the following two bundles, where each item is represented by its level for some specific agent:

$$X = \{8, 4, 2\} \quad Y = \{7, 6\}$$

Note that $|X| > |Y|$, X is lexicographically-better than Y , and even the Borda score of X is higher. However, the level of the two best items of X is only 12, which is less than that of the best two items of Y . Hence, by Theorem 1, it is not true that $X \succeq_i^{NDD} Y$. Indeed, X is not better than Y according to all DD valuations. For example, let $u(x) = \text{LEV}(x)^2$. The function u has DD, but $u(X) = 82 < 83 = u(Y)$. \square

Corollary 1. *There is a linear-time algorithm to check whether $X \succeq_i^{NDD} Y$*

The algorithm is obvious from Theorem 1 and is omitted.

Theorem 2. *Let X, Y be two multi-bundles. Then $X \succeq_i^{PDD} Y$ if and only if $Y \not\succeq_i^{NDD} X$ or $\text{LEV}_i(Y) = \text{LEV}_i(X)$.*

Proof. \Leftarrow : If $Y \not\succeq_i^{NDD} X$, then by definition of NDD, X has strictly more utility than Y for some consistent DD utility function. If $\text{LEV}_i(X) = \text{LEV}_i(Y)$, then X has the same Borda utility as Y . In both cases $X \succeq_i^{PDD} Y$.

\Rightarrow : Suppose that $Y \succeq_i^{NDD} X$ and $\text{LEV}_i(Y) \neq \text{LEV}_i(X)$. So $\text{LEV}_i(Y) > \text{LEV}_i(X)$. By Theorem 1 we must have $|Y| \geq |X|$, and for every $k \leq |X|$: $\text{LEV}_i(Y_k) \geq \text{LEV}_i(X_k)$.

For every k , Theorem 1 is applicable to the multi-bundles X_k, Y_k . So for any function $u_i \in \mathcal{U}^{DD}(\succ_i)$, $u_i(Y_k) \geq u_i(X_k)$. Moreover, if $\text{LEV}_i(Y_{|X|}) > \text{LEV}_i(X_{|X|})$ then also $u_i(Y_{|X|}) > u_i(X_{|X|}) = u_i(X)$; otherwise, $\text{LEV}_i(Y_{|X|}) = \text{LEV}_i(X_{|X|})$ and $u_i(Y_{|X|}) = u_i(X_{|X|})$, so Y must have more than $|X|$ items, and since all items have positive utility, again $u_i(Y) > u_i(X)$. In all cases, for all DD consistent utility functions, $u_i(Y) > u_i(X)$. Hence, $X \not\succeq_i^{PDD} Y$. \square

Corollary 2. *There is a linear-time algorithm to check whether $X \succeq_i^{PDD} Y$.*

Corollary 3. *It can be decided in polynomial time whether a given allocation is NDDPR, NDDEF, PDDPR or PDDEF.*

5 Existence of NDD-Proportional Allocations

In this section, we prove the following necessary and sufficient condition for existence of NDDPR allocations.

Theorem 3. *An NDDPR allocation exists if-and-only-if:*

- (a) *The number of items is $M = m \cdot n$, where m is an integer and n is the number of agents, and -*
- (b) *Each agent has a different best item.*

In case it exists, it can be found in time $O(M)$.

Proof of \Rightarrow . Suppose the two conditions are satisfied. We are going to prove that the following simple algorithm produces an NDDPR allocation:

Repeat as long as there are items:
 For $i = 1 \dots n$: give agent i his best remaining item.
 For $i = n \dots 1$: give agent i his best remaining item.

The proof proceeds in two steps. First, we identify a “worst case” scenario of the algorithm. Second, we prove that even in this scenario, the allocation is NDDPR.

We will prove that the worst-case scenario is when (1) all agents have the same set of n best items and the same set of $M - n$ worst items, and (2) all agents rank the $M - n$ worst items in exactly the same way (note that the agents must differ in their ranking on the set of n best items since each agent must still have a different top item).

For each agent i , we define two bundles:

X_i — the bundle allocated to i by the above algorithm.

Y_i — the bundle that would be allocated to i in the “worst-case” scenario.

We will prove that the following are true for each agent i :

- (a) $X_i \succeq_i^{Nec} Y_i$ (i.e, Y_i is indeed the worst-case outcome of the algorithm),
- (b) $n \cdot Y_i \succeq_i^{NDD} \mathcal{M}$ (even this worst-case is NDDPR).

The theorem follows from (a) and (b) by transitivity.

To prove (a), we use the following well-known characterization [Brams *et al.*, 2013; Bouveret *et al.*, 2010]:

$X \succeq_i^{Nec} Y$ if-and-only-if there exists a bijection $f : X \rightarrow Y$ such that for all $x \in X$: $x \succeq f(x)$.

We use the bijection that maps the best item in X_i to the best item in Y_i , the 2nd-best item in X_i to the 2nd-best item in Y_i , and so on. We prove that for every k , the k -th best item in X_i is at least as good as the k -th best item in Y_i . This is certainly true for $k = 1$ since the best items are the same. For $k > 1$, when agent i picks his k -th item (which is the k -th best item in X_i), the number of items already taken is $kn - i$ (if k is even) or $kn - (n - i + 1)$ (if k is odd). Therefore, agent i 's best remaining item has a level of at least $M - (kn - i)$ (if k is even) or $M - kn + (n - i + 1)$ (if k is odd). This is always at least as large as the level of the k -th-best item in Y_i .

To prove (b), we apply Theorem 1. The first condition is $|n \cdot Y_i| \geq |\mathcal{M}|$. Indeed, in our case $|n \cdot Y_i| = n \cdot m = M = |\mathcal{M}|$. The second condition is that, for each integer $k' \in \{1, \dots, M\}$, the level of the best k' items in $n \cdot Y_i$ minus the level of the best k' items in \mathcal{M} is weakly positive. The following table shows the levels and their differences for $k' \in \{1, \dots, n\}$ (recall that all items in \mathcal{M} are distinct):

$n \cdot Y_i$	M	M	\dots	M
\mathcal{M}	M	$M - 1$	\dots	$M - n + 1$
Difference	0	1	\dots	$n - 1$

Clearly the level difference is always weakly positive. The total level difference after the best n items is $n(n - 1)/2$.

The following table shows the levels and their differences for $k' \in \{n + 1, \dots, 2n\}$:

$n \cdot Y_i$	$M - 2n + i$	$M - 2n + i$	\dots	$M - 2n + i$
\mathcal{M}	$M - n$	$M - n - 1$	\dots	$M - 2n + 1$
Difference	$i - n$	$i - n + 1$	\dots	$i - 1$

Clearly, since $i \geq 1$, the level difference in the table is weakly larger than $(1 - n) + (1 - n + 1) + \dots + (1 - 1) = -n(n - 1)/2$. Since we already have a level difference of $+n(n - 1)/2$ from the previous n items, the total level difference remains weakly positive. The total level difference in the table is $-n(n - 1)/2 + n(i - 1)$ so the total level difference of the best $2n$ items is $n(i - 1)$.

The following table shows the levels and their differences for $k' \in \{2n + 1, \dots, 3n\}$:

$n \cdot Y_i$	$M - 2n - i + 1$	$M - 2n - i + 1$...	$M - 2n - i + 1$
\mathcal{M}	$M - 2n$	$M - 2n - 1$...	$M - 3n + 1$
Difference	$1 - i$	$2 - i$...	$n - i$

Clearly, each difference in the table is $-(i - 1)$ or larger, and there are n such differences. Since we already have a level difference of $+n \cdot (i - 1)$ from the previous $2n$ items, the total level difference remains weakly positive. The total level difference in the table is $n(n - 1)/2 - (n - 1)i$ so the total level difference of the best $3n$ items is $n(n - 1)/2$.

We are now back at the same situation as after the best n items, so we can repeat the same reasoning and conclude that the level difference of the best k' items is always weakly-positive. \square

Proof of \Leftarrow . Let $(p(1), \dots, p(n))$ be an NDDPR allocation. Then for every agent i , $n \cdot X_i \succ_i^{NDD} \mathcal{M}$. We now apply the two conditions of Theorem 1.

(a) For all i : $|n \cdot X_i| \geq |\mathcal{M}| \implies n \cdot |X_i| \geq M$. But this must be an equality since the total number of items in all n bundles is exactly M . Therefore, the total number of items is $n \cdot |X_i|$ which is an integer multiple of n .

(b) For all i , the level of the best item in $n \cdot X_i$ must be weakly larger than the level of the best item in \mathcal{M} . So for every i , X_i must contain agent i 's best item. So the best items of all agents must be different. \square

6 NDD-Proportionality in Simulations

We compared the various fairness criteria using a quick preliminary simulation experiment with 2 agents. First, we checked to what extent the probability that an NDDPR allocation exists is higher than the probability that an NecPR allocation exists. These probabilities naturally depend on the correlation between the agents' rankings. When the rankings are completely correlated, both NDDPR and NecPR allocations do not exist; when the rankings are completely independent, both NDDPR and NecPR allocations exist with high probability; the interesting zone is when the rankings are partially correlated. To simulate such rankings, we determined for each item a "market value" drawn uniformly at random from $[1, 2]$. We determined the value of each item to each agent as the item's market value plus noise drawn uniformly at random from $[-A, A]$, where A is a parameter. Based on the cardinal values we determined the agent's ordinal ranking. Then, we checked whether there exists an NecPR/NDDPR allocation. We did this experiment 1000 times for different values of $A \in \{0.1, \dots, 1\}$ and for different numbers of items — $2m$ items for $m \in \{2, \dots, 8\}$. Typical results are plotted in Figure 1; the probability of existence of NDDPR allocations (balls) is clearly higher than that of NecPR allocations (triangles):

Since we had randomly-generated cardinal values, we used them for a secondary purpose — we checked whether, when an NDDPR allocation exists, the one found by the simple algorithm of Section 5 is proportional according to these values (dashed lines). Note that, since the randomization we

used is completely uniform and does not use the DD assumption, the probability that DD holds for both agents is very low — $1/((2m - 1)!)^2$. Nevertheless, our NDDPR allocation (when it exists) is almost always proportional when the number of items or the noise size are sufficiently large, which further shows the robustness of our algorithm.

We also checked the probability of existence of PDDPR and PosPR allocations, and it was nearly 1.0. Thus, apparently the NecPR requirement is too strong and the PDDPR and PosPR requirements are too weak, while the NDDPR requirement hits a sweet spot between 'recall' and 'precision': it allows us to solve many instances (= high 'recall') and most solutions are satisfactory (= high 'precision').

7 Pareto-efficiency and Envy-freeness

While the focus of this paper is on the effect of the DD condition on proportionality, we briefly survey below some results that we have for Pareto-efficiency and envy-freeness.

7.1 Pareto-efficiency

In the preceding sections we saw that the DD condition has a substantial effect on fairness notions: NDD-fair allocations are strictly easier (in terms of existence) than necessary-fair allocations, and PDD-fair allocations are harder than possibly-fair allocations (see examples in Subsection 3.2). Interestingly, the DD condition does not have this effect on Pareto-efficiency (PE).

An allocation is called *Necessarily (NDD) PE* if it is PE according to all profiles with additive (DD) utility functions.

Theorem 4. *Necessary-PE is equivalent to NDD-PE.*

Proof. \implies : If an allocation is necessarily-PE, then by definition it is also NDD-PE.

\Leftarrow : Suppose an allocation X is not necessarily-PE. Then, by Aziz *et al.* [2016b] Thm 9, there are two options:

- (i) X is not possibly-PE. Then, it is certainly not NDD-PE.
- (ii) X admits a Pareto-improving one-for-two-swap. I.e., there are two agents Alice and Bob, such that X_{Bob} contains an item which Alice strictly prefers over two items in X_{Alice} . Then X is not NDD-PE, since it is not PE according to the utility profile in which u_{Alice} is lexicographic, and $u_{Bob}(x) = M^2 + \text{LEV}_{Bob}(x)$, so that his utility is dominated by the number of items he has. \square

Analogously, Possible-DD-PE is equivalent to Possible-PE; proof is omitted. It is interesting that whereas DD leads to new fairness notions, it does not lead a new efficiency notion.

7.2 DD-envy-freeness

Since every NDDEF allocation is NDDPR, the two conditions of Theorem 3 are necessary for the existence of NDDEF allocations for any number of agents. In the special case of $n = 2$ agents, NDDPR is equivalent to NDDEF so these conditions are also sufficient. But for $n = 3$ they are no longer sufficient.

Example 4. *There are six items $\{1, \dots, 6\}$. The preferences of the three agents Alice Bob and Carl are:*

$$\begin{aligned} \text{Alice: } & 6 \succ 5 \succ 3 \succ 4 \succ 2 \succ 1 \\ \text{Bob: } & 5 \succ 4 \succ 3 \succ 6 \succ 2 \succ 1 \end{aligned}$$




Figure 1: NDD-proportionality in Simulations

Carl: $4 \succ 6 \succ 3 \succ 5 \succ 2 \succ 1$

The conditions of Theorem 3 are clearly satisfied. However, no NDDEF allocation exists. Proof: The preferences are the same up to a cyclic permutation between 6 5 and 4, so the agents are symmetric and it is without loss of generality to assume that Alice receives item 1. Therefore Alice’s bundle is $\{6, 1\}$ and her Borda score is 7. To ensure that Alice is not envious, Bob must get $\{5, 2\}$ and Carl must get $\{3, 4\}$. This allocation is NDDPR but it is not NDDEF, since Bob Borda-envies Carl. \square

Since the NDDPR characterization does not work for NDDEF allocations even for three agents, it is an open problem whether the existence of NDDEF allocations can be decided efficiently. When the number of agents is not bounded, we have the following hardness result:

Theorem 5. *Checking the existence of NDDEF allocations is NP-complete when there are n agents and at least $2n$ items.*

The reduction is similar to the NP-completeness of checking existence of NEF allocations [Bouveret *et al.*, 2010]. The proof requires carefully checking that the reduction argument works for NDDEF as well.

When the number of agents is *constant* (at least 3) and the number of items is variable, the runtime complexity of checking NDDEF existence is an open question: is it polynomial like NDDPR, or NP-hard like NEF [Aziz *et al.*, 2016c]?

8 Conclusions and Future Work

We formalized natural ways to compare sets by using the DD (diminishing differences) assumption. The relations lead to new fairness concepts which we studied in detail. In future work, it will be interesting to identify other interesting set extensions that correspond to classes of utility functions. Other future work includes extending some of our results to the case where agents may express weak preferences, or where the items have negative utilities (chores).

Acknowledgments

We acknowledge the Dagstuhl Seminar 16232 on Fair Division where this project was initiated. We are grateful to four

anonymous IJCAI reviewers for their very helpful comments.

Haris Aziz is supported by a Julius Career Award. Erel is supported by the ISF grant 1083/13, the Doctoral Fellowships of Excellence Program and the Mordecai and Monique Katz Graduate Fellowship Program at Bar-Ilan University. Avinatan Hassidim is supported by ISF grant 1394/16.

References

[Abdulkadiroglu and Sönmez, 2003] Atila Abdulkadiroglu and Tayfun Sönmez. School choice: A mechanism design approach. *The American Economic Review*, 93(3):729–747, 2003.

[Ashlagi *et al.*, 2014] Itai Ashlagi, Mark Braverman, and Avinatan Hassidim. Stability in large matching markets with complementarities. *Operations Research*, 62(4):713–732, 2014.

[Aziz *et al.*, 2015] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, October 2015.

[Aziz *et al.*, 2016a] H. Aziz, J. Lang, and J. Monnot. Computing Pareto Optimal Committees. In *Proc. IJCAI-16*, pages 60–66, 2016.

[Aziz *et al.*, 2016b] Haris Aziz, Páter Biró, Jérôme Lang, Julien Lesca, and Jérôme Monnot. Optimal Reallocation Under Additive and Ordinal Preferences. In *Proc. AAMAS-16*, pages 402–410, 2016.

[Aziz *et al.*, 2016c] Haris Aziz, Ildiko Schlotter, and Toby Walsh. Control of Fair Division. In *Proc. IJCAI-16*, pages 67–73, 2016.

[Aziz *et al.*, 2017] Haris Aziz, Gerhard Rauchecker, Guido Schryen, and Toby Walsh. Algorithms for Max-Min Share Fair Allocation of Indivisible Chores. In *Proc. AAI-17*, pages 1–7, 2017.

[Aziz, 2015] Haris Aziz. A note on the undercut procedure. *Social Choice and Welfare*, 45(4):723–728, 2015.

- [Barberà *et al.*, 2001] S. Barberà, B. Dutta, and A. Sen. Strategy-proof social choice correspondences. *Journal of Economic Theory*, 101(2):374–394, 2001.
- [Barberà *et al.*, 2004] S. Barberà, W. Bossert, and P. K. Pattanaik. Ranking sets of objects. In S. Barberà, P. J. Hammond, and C. Seidl, editors, *Handbook of Utility Theory*, volume II, chapter 17, pages 893–977. Kluwer Academic Publishers, 2004.
- [Baumeister *et al.*, 2016] Dorothea Baumeister, Sylvain Bouveret, Jérôme Lang, Nhan-Tam Nguyen, TrungThanh Nguyen, Jörg Rothe, and Abdallah Saffidine. Positional scoring-based allocation of indivisible goods. *Autonomous Agents and Multi-Agent Systems*, pages 1–28, 2016.
- [Bouveret and Lang, 2011] Sylvain Bouveret and Jérôme Lang. A General Elicitation-free Protocol for Allocating Indivisible Goods. In *Proc. IJCAI-11*, pages 73–78. AAAI Press, 2011.
- [Bouveret *et al.*, 2010] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Fair Division Under Ordinal Preferences: Computing Envy-Free Allocations of Indivisible Goods. In *Proc. ECAI-10*, pages 387–392, Amsterdam, The Netherlands, The Netherlands, 2010. IOS Press.
- [Brams and Kaplan, 2004] Steven J. Brams and Todd R. Kaplan. Dividing the Indivisible. *Journal of Theoretical Politics*, 16(2):143–173, 2004.
- [Brams and Taylor, 2000] Steven J. Brams and Alan D. Taylor. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody (Norton Paperback)*. W. W. Norton & Company, reprint edition, October 2000.
- [Brams *et al.*, 2012] Steven J. Brams, Kilgour, and Christian Klamler. The undercut procedure: an algorithm for the envy-free division of indivisible items. *Social Choice and Welfare*, 39(2-3):615–631, October 2012.
- [Brams *et al.*, 2013] Steven J. Brams, D. Marc Kilgour, and Christian Klamler. Two-Person Fair Division of Indivisible Items: An Efficient, Envy-Free Algorithm. *Social Science Research Network Working Paper Series*, pages 130–141, June 2013.
- [Brandt, 2017] F. Brandt. Rolling the dice: Recent results in probabilistic social choice. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 1. AI Access, 2017.
- [Bronfman *et al.*, 2015a] Slava Bronfman, Noga Alon, Avinatan Hassidim, and Assaf Romm. Redesigning the israeli medical internship match. In *Proc. EC-15*, pages 753–754. ACM, 2015.
- [Bronfman *et al.*, 2015b] Slava Bronfman, Avinatan Hassidim, Arnon Afek, Assaf Romm, Rony Shreberk, Ayal Hassidim, and Anda Massler. Assigning israeli medical graduates to internships. *Israel journal of health policy research*, 4(1):6, 2015.
- [Budish, 2011] Eric Budish. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6):1061–1103, December 2011.
- [Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Herve Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The Unreasonable Fairness of Maximum Nash Welfare. In *Proc. EC-16*, 2016.
- [Cho, 2012] Wonki J. Cho. Probabilistic assignment: A two-fold axiomatic approach. Job-Market Paper, 2012.
- [Darmann and Klamler, 2016] Andreas Darmann and Christian Klamler. Proportional Borda allocations. *Social Choice and Welfare*, 47(3):543–558, 2016.
- [Elkind and Lackner, 2014] Edith Elkind and Martin Lackner. On Detecting Nearly Structured Preference Profiles. In *Proc. AAI-14*, pages 661–667. AAAI Press, 2014.
- [Elkind *et al.*, 2017] E. Elkind, M. Lackner, and D. Peters. Structured preferences. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 10. 2017.
- [Hadar and Russell, 1969] Joseg Hadar and William R. Russell. Rules for ordering uncertain prospects. *American Economic Review*, 59:2–34, 1969.
- [Hassidim *et al.*, 2016a] Avinatan Hassidim, Assaf Romm, and Ran I Shorrer. Strategic behavior in a strategy-proof environment. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 763–764. ACM, 2016.
- [Hassidim *et al.*, 2016b] Avinatan Hassidim, Assaf Romm, Ran I Shorrer, et al. Redesigning the israeli psychology masters match. Technical report, 2016.
- [Herreiner and Puppe, 2002] Dorothea Herreiner and Clemens Puppe. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19(2):415–430, 2002.
- [Kalinowski *et al.*, 2013] Thomas Kalinowski, Nina Narodytska, and Toby Walsh. A Social Welfare Optimal Sequential Allocation Procedure. In *Proc. IJCAI-13*, pages 227–233. AAAI Press, 2013.
- [Mahajne *et al.*, 2015] Muhammad Mahajne, Shmuel Nitzan, and Oscar Volij. Level r consensus and stable social choice. *Social Choice and Welfare*, 45(4):805–817, 2015.
- [Nguyen *et al.*, 2015] N-T. Nguyen, D. Baumeister, and J. Rothe. Strategy-proofness of scoring allocation correspondences for indivisible goods. In *Proc. IJCAI-15*, pages 1127–1133. AAAI Press, 2015.
- [Nitzan *et al.*, 2017] Mor Nitzan, Shmuel Nitzan, and Erel Segal-Halevi. On Level-1 Consensus Ensuring Stable Social Choice, April 2017. arXiv preprint 1704.06037.
- [Procaccia and Wang, 2014] Ariel D. Procaccia and Junxing Wang. Fair Enough: Guaranteeing Approximate Maximin Shares. In *Proc. EC-14*, pages 675–692, New York, NY, USA, 2014. ACM.
- [Roth and Peranson, 1997] A. E. Roth and E. Peranson. The effects of the change in the NRMP matching algorithm. *JAMA*, 278(9):729–732, 1997.
- [Young, 1974] H.P. Young. An axiomatization of borda’s rule. *Journal of Economic Theory*, 9(1):43 – 52, 1974.