Two Desirable Fairness Concepts for Allocation of Indivisible Objects under Ordinal Preferences

HARIS AZIZ
and
SERGE GASPERS
and
SIMON MACKENZIE
and
TOBY WALSH
Data61 and UNSW

Fair allocation of indivisible objects under ordinal preferences is an important problem. Unfortunately, a fairness notion like envy- freeness is both incompatible with Pareto optimality and is also NP-complete to achieve. To tackle this predicament, we consider a different notion of fairness, namely proportionality. We frame allocation of indivisible objects as randomized assignment but with integrality requirements. We then use the stochastic dominance relation to define two natural notions of proportionality. Since an assignment may not exist even for the weaker notion of proportionality, we propose relaxations of the concepts — optimal weak proportionality and optimal proportionality. For both concepts, we propose algorithms to compute fair assignments under ordinal preferences. Both new fairness concepts appear to be desirable in view of the following: they are compatible with Pareto optimality, admit efficient algorithms to compute them, are based on proportionality, and are guaranteed to exist.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems; I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Theory, Algorithms, Economics

Additional Key Words and Phrases: Fair allocation, Multiagent resource allocation

1. INTRODUCTION

The principled allocation of resources is one of most pressing problems faced by society [see, e.g., Bezáková and Dani, 2005, Bouveret and Lemaître, 2014]. Within the rich field of resource allocation, a typical allocation setting has a set of agents $N = \{1, \dots, n\}$, a set of objects $O = \{o_1, \dots, o_m\}$ with $m \geq n$ and each agent $i \in N$ expressing complete and transitive *ordinal* preferences $\succsim_i$ over O. The goal is to allocate *all* the objects in O to the agents in a fair manner. Since eliciting preferences over bundles of objects requires exponential time, we only assume that agents express preferences over individual objects. These preferences over objects can involve indifference ($\succ_i$ denotes strict preference whereas $\sim_i$ denotes indifference). The setting is referred to as the *assignment problem* or the *house allocation problem* [see, e.g., Baumeister et al., 2014, Bouveret et al., 2010, Gärdenfors, 1973,

Authors' addresses: {`haris.aziz,serge.gaspers,simon.mackenzie,toby.walsh`}`@nicta.com.au`

Manlove, 2013, Pruhs and Woeginger, 2012, Wilson, 1977]. The model is applicable to many resource allocation or fair division settings in which multiple objects may be allocated to agents.

Example 1 Assignment problem.

$$1: \quad o_1 \succ_1 o_2 \succ_1 o_3 \succ_1 o_4 \succ_1 o_5 \qquad 2: \quad o_2 \sim_2 o_3 \succ_2 o_1 \sim_2 o_4 \sim_2 o_5$$

Since the goal is to identify fair allocations, there is a need to formalize what fairness entails. Two of the most fundamental concepts of fairness are *envy-freeness* and *proportionality*. Envy-freeness requires that no agent prefers another agent's allocation. Proportionality requires that each agent should get an allocation that gives him at least $1/n$ of the utility that he would get if he got all the objects. When agents' ordinal preferences are known but utility functions are not given, then ordinal notions of envy-freeness and proportionality need to be formulated. Since envy-freeness is defined by comparing agents' allocations, in order to reason about envy-freeness we need to make some assumptions on how preferences over objects are extended to preferences over objects. One basic assumption we can make is that of *responsiveness*: if an agent gets an extra object or one of his objects is replaced by a strictly more preferred object, he is happier. Based on the assumption of responsiveness, one can define weak envy-freeness (another agent's allocation is not strictly more preferred) and strong envy-freeness (one's allocation is weakly preferred over others' allocations). Not only are both notions incompatible with Pareto optimality but it is NP-complete to check whether an envy-free assignment exists [Aziz et al., 2014, Bouveret et al., 2010]. In view of this, we seek a fairness concept that satisfies the following requirements: (1) captures a natural fairness goal along similar lines as envy-freeness or proportionality, (2) guaranteed to exist, (3) efficiently computable[1], and (4) compatible with Pareto optimality.

2. PROPORTIONALITY BASED ON STOCHASTIC DOMINANCE

We take proportionality as a starting point and can define an ordinal version of proportionality. On face value, proportionality appears to be based on cardinal utilities. Indeed one can superimpose cardinal utilities consistent with the ordinal preferences and check whether a proportional assignment exists. However checking whether a proportional assignment exists for cardinal utilities is also NP-complete [Demko and Hill, 1988]. We adopt instead a different perspective in which we define ordinal notions of proportionality by first considering discrete assignments as special kind of random assignments in which each agent gets a fraction (probability) of getting each object as follows. We will still compute fair discrete assignments but will do so via comparisons with fractional assignments.

A fractional assignment p is a $(n \times m)$ matrix $[p(i)(o_j)]$ such that $p(i)(o_j) \in [0,1]$ for all $i \in N$, and $o_j \in O$, and $\sum_{i \in N} p(i)(o_j) = 1$ for all $j \in \{1, \ldots, m\}$. The value $p(i)(o_j)$ represents the probability of object o_j being allocated to agent i. Each row $p(i) = (p(i)(o_1), \ldots, p(i)(o_m))$ represents the allocation of agent i. The columns correspond to the objects $o_1, \ldots, o_m$. A fractional assignment is *discrete* if $p(i)(o) \in \{0,1\}$ for all $i \in N$ and $o \in O$.

[1] The input is agents expressing preferences over objects. Hence we will say that an algorithm is polynomial-time if it runs in time polynomial in the number of agents and objects.

We then use the stochastic dominance relation [see e.g., Aziz et al., 2013] to compare fractional allocations. Informally, an agent 'SD-prefers' one allocation over another if for each object o, the former allocation gives the agent as many units of objects that are at least preferred as o as the latter allocation. More formally, given two fractional assignments p and q, $p(i) \succsim_i^{SD} q(i)$, i.e., agent i *SD prefers* allocation $p(i)$ to allocation $q(i)$ if

$$\sum_{o_j \in \{o_k : o_k \succsim_i o\}} p(i)(o_j) \geq \sum_{o_j \in \{o_k : o_k \succsim_i o\}} q(i)(o_j) \text{ for all } o \in O.$$

Agent i *strictly SD prefers* $p(i)$ to $q(i)$ if $p(i) \succsim_i^{SD} q(i)$ and $\neg[q(i) \succsim_i^{SD} p(i)]$. SD can also be viewed from a utility perspective which underlines its fundamental nature: an agent prefers one allocation over another with respect to the SD relation if he gets at least as much utility from the former allocation as the latter for all cardinal utilities consistent with the ordinal preferences.

Based on SD, one can define two fairness notions [Aziz et al., 2014]. In particular, we define *weak SD proportionality* as requiring that no agent strictly prefers the allocation in which $1/n$ of each object is obtained to his own allocation. We define *SD proportionality* as requiring that each agent weakly SD-prefers his allocation over the allocation in which $1/n$ of each object is obtained. An assignment p satisfies *weak SD proportionality* if no agent strictly SD prefers the uniform assignment to his allocation: $\neg[(1/n, \dots, 1/n) \succ_i^{SD} p(i)]$ for all $i \in N$. An assignment p satisfies *SD proportionality* if each agent SD prefers his allocation to the allocation under the uniform assignment: $p(i) \succsim_i^{SD} (1/n, \dots, 1/n)$ for all $i \in N$.

SD proportionality and weak SD proportionality are not only desirable fairness concepts but they are also computationally more tractable than ordinal notions of envy-freeness.

THEOREM 1 [AZIZ ET AL., 2014]. *We can check in polynomial time whether a discrete SD proportional assignment exists even if agents are allowed to express indifference between objects. For a constant number of agents, we can check in polynomial time whether a weak SD proportional discrete assignment exists.*

A possible criticism of weak SD proportionality and SD proportionality concepts is that even the weaker of the two is not achievable in general. Consider the following example. The weak SD proportionality constraint is violated for the agent who gets at most one object.

EXAMPLE 2. *Assume that the preferences of the agents are as follows.*

$1: \quad o_1 \sim_1 o_2 \sim_1 o_3 \qquad\qquad 2: \quad o_1 \sim_2 o_2 \sim_2 o_3$

3. OPTIMAL PROPORTIONALITY AND OPTIMAL WEAK PROPORTIONALITY

When a weak SD proportional assignment does not exist, we would still like to allocate the objects in a principled manner. We relax weak SD proportionality and SD proportionality to propose optimal proportionality and optimal weak proportionality.

Definition 1 *Optimal proportionality [Aziz et al., 2015].* We say that an assignment satisfies $1/\alpha$ *proportionality* if $p(i) \succsim_i^{SD} (1/\alpha, \dots, 1/\alpha)$ for all $i \in N$. We note that $1/n$- proportionality is equivalent to SD proportionality. An assignment satisfies *optimal proportionality* if $p(i) \succsim_i^{SD} (1/\alpha, \dots, 1/\alpha)$ for all $i \in N$ for the smallest possible α. We will refer to the smallest such α as α^* and call $1/\alpha^*$ as the *optimal proportionality value.*

Definition 2 *Optimal weak proportionality [Aziz et al., 2015].* Just like the concept of SD proportionality can be used to define optimal proportionality, weak SD proportionality can be used to define *optimal weak proportionality.* We say that an assignment satisfies $1/\beta$ *weak proportionality* if $(1/\beta, \dots, 1/\beta) \not\succ_i^{SD} p(i)$ for all $i \in N$. We note that $1/n$ weak proportionality is equivalent to weak SD proportionality. An assignment satisfies *optimal weak proportionality* if $(1/\beta, \dots, 1/\beta) \not\succ_i^{SD} p(i)$ for all $i \in N$ for the infimum of the set $\{\beta \mid \exists \text{ a } 1/\beta \text{ weak proportional assignment}\}$. We will refer to the infimum as β^* and call $1/\beta^*$ as the *optimal weak proportionality value.*

Theorem 1 can be generalized from $1/n$ proportionality to $1/\alpha$ proportionality for any value of α. The algorithm can be used to check the existence of a $1/\alpha$ proportional assignment for different values of α. However, among other cases, if $m < n$, we then know that a $1/\alpha$ proportional assignment does not exist for any finite value of α. We show that α^* is finite if and only if there exists an assignment in which each agent gets one of his most preferred objects. Since α is a positive real in the interval $(0, \infty]$, it may appear that even binary search cannot be used to find the optimal proportional assignment in polynomial time. Interestingly, we only need to check a polynomial number of values of α to find the optimal proportional assignment.

Theorem 2 [Aziz et al., 2015]. *An optimal proportional assignment can be computed in polynomial time.*

We point out that an SD proportional assignment (if it exists) is an optimal proportional assignment. Moreover, even if an SD proportional assignment does not exist, an optimal proportional assignment suggests a desirable allocation of objects. For example, for the preference profile in Example 2, we observed that there exists no weak SD proportional assignment. On the other hand, the assignment that gives two objects to one agent and one object to the other is an optimal proportional assignment where the optimal proportionality value is 1/3.

In a similar approach as for optimal proportionality, for a constant number of agents, it can be checked in polynomial time whether a $1/\beta$ weak proportional discrete assignment exists. We also show that for any assignment setting, $\beta^* \geq 1$ and is finite if and only if $m \geq n$.

Theorem 3 [Aziz et al., 2015]. *If the number of agents is constant, an optimal weak proportional assignment can be computed in polynomial time.*

We note that whereas an SD proportional assignment is an optimal proportional assignment, a weak SD proportional assignment may not be an optimal weak proportional assignment.

EXAMPLE 3. *Assume that the preferences of the agents are as follows.*

$$1: \quad o_1 \succ_1 o_2 \succ_1 o_3 \succ_1 o_4 \succ_1 o_5$$
$$2: \quad o_2 \sim_2 o_3 \succ_2 o_1 \sim_2 o_4 \sim_2 o_5$$

Note that the assignment p that gives $\{o_2, o_3\}$ to agent 1 and the other objects to agent 2 is weak SD proportional. In fact it is not only $1/2$ weak proportional but $(3/5 - \epsilon))$ weak proportional where $\epsilon > 0$ is arbitrarily small. It is not $1/\beta$ weak proportional for $1/\beta < 3/5$. We now consider an assignment q, that gives $\{o_1\}$ to agent 1 and the other objects to agent 2. But q is $(1 - \epsilon)$ weak proportional where $\epsilon > 0$ is arbitrarily small. This shows that a weak SD proportional discrete assignment may not be an optimal weak proportional assignment.

4. CONCLUSIONS AND OPEN PROBLEMS

In this note, we highlighted two desirable fairness concepts that have recently been proposed [Aziz et al., 2015]. The most interesting remaining problem is checking whether there exists a polynomial-time algorithm for computing an optimal weak proportional assignment when the number of agents is not constant.

Acknowledgments

We thank the reviewers of the AAMAS 2014 conference and the Artificial Intelligence journal for their helpful reviews. The work is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. Serge Gaspers is the recipient of an Australian Research Council Discovery Early Career Researcher Award (project number DE120101761) and a Future Fellowship (project number FT140100048).

REFERENCES

H. Aziz, F. Brandt, and M. Brill. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In *Proc. of 12th AAMAS Conference*, pages 455–462. IFAAMAS, 2013.

H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. In *Proc. of 13th AAMAS Conference*, pages 1305–1312, 2014.

H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.

D. Baumeister, S. Bouveret, J. Lang, N-T. Nguyen, T-T. Nguyen, J. Rothe, and A. Saffidine. Axiomatic and computational aspects of scoring allocation rules for indivisible goods. In *Proc. of 5th COMSOC*, 2014.

I. Bezáková and V. Dani. Allocating indivisible goods. *SIGecom Exchanges*, 5(3):11–18, 2005.

S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proc. of 13th AAMAS Conference*, pages 1321–1328. IFAAMAS, 2014.

S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proc. of 19th ECAI*, pages 387–392, 2010.

S. Demko and T. P. Hill. Equitable distribution of indivisible objects. *Mathematical Social Sciences*, 16:145–158, 1988.

P. Gärdenfors. Assignment problem based on ordinal preferences. *Management Science*, 20:331–340, 1973.

D. Manlove. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company, 2013.

K. Pruhs and G. J. Woeginger. Divorcing made easy. In *Proceedings of FUN*, number 7288 in LNCS, pages 305–314, 2012.

L. Wilson. Assignment using choice lists. *Operations Research Quarterly*, 28(3):569—578, 1977.