Econ Theory Bull (2016) 4:307–324
DOI 10.1007/s40505-015-0089-1

CrossMark

RESEARCH ARTICLE

A generalization of the AL method for fair allocation of indivisible objects

Haris Aziz[1]

Received: 14 December 2014 / Accepted: 28 November 2015 / Published online: 23 December 2015

Abstract We consider the assignment problem in which agents express ordinal preferences over m objects and the objects are allocated to the agents based on the preferences. In a recent paper Brams, Kilgour, and Klamler (Not AMS 61(2):130–141, 2014) , presented the AL method to compute an envy-free assignment for two agents. The AL method crucially depends on the assumption that agents have strict preferences over objects. We generalize the AL method to the case where agents may express indifferences and prove the axiomatic properties satisfied by the algorithm. As a result of the generalization, we also get a $O(m)$ speedup on previous algorithms to check whether a complete envy-free assignment exists or not. Finally, we show that unless $P = NP$, there can be no polynomial time extension of GAL to the case of arbitrary number of agents.

Keywords Fair division · Envy-freeness · Pareto optimality · AL method

JEL Classification C70 · D61 · D71

1 Introduction

Fair allocation of resources is one of the most critical issues faced by society. A basic, yet widely applicable, problem in computer science and economics is to allocate discrete objects to agents given the ordinal preferences of the agents over the objects. The setting is referred to as the *assignment problem* or the *house allocation problem* (see, e.g., Abraham et al. 2005; Aziz et al. 2014; Bouveret et al. 2010; Brams and Kaplan 2004; Brams et al. 2003; Brams and Fishburn 2000; Brams et al. 2012; Demko and Hill

✉ Haris Aziz
haris.aziz@nicta.com.au

1 NICTA and UNSW, Sydney 2033, Australia

1988; Gärdenfors 1973; Manlove 2013; Wilson 1977; Young 1995). In this setting, there are a set of agents $N = \{1, \ldots, n\}$ and a set of objects $O = \{o_1, \ldots, o_m\}$ with each agent $i \in N$ expressing ordinal preferences $\succsim_i$ over O. Each object is assumed to be acceptable to the agents. The goal is to allocate the objects among the agents in a fair or optimal manner without allowing transfer of money. The model is applicable to many resource allocations or fair division settings, where the objects may be public houses, school seats, course enrolments, kidneys for transplant, car park spaces, chores, joint assets of a divorcing couple, or time slots in schedules.

For the assignment problem, the case of two agents is especially central. Many disputes are between two parties and may require division of common resources. Divorce proceedings are one of the settings in which common assets need to be divided among the two parties. Other examples in history include partition of countries which results in the need to divide common assets.

When objects are allocated among agents, it is desirable that they are allocated in a fair and efficient manner. For fairness, one of the most established concepts is *envy-freeness*. A formal study of envy-freeness in microeconomics can be traced back to the work of Foley (1967). Envy-freeness requires that each agent should prefer its allocation over other agents' allocations. Envy-freeness can be trivially satisfied by not giving any objects to any agents. However, if we insist that the assignment should be *complete*, i.e., it allocated all the objects to the agents, no assignment may be envy-free as is the case in which there is only one object and the agent who does not get any object is envious. The most established notion of efficiency is Pareto optimality which requires that there should be no other allocation which each agent weakly prefers and at least one agent strictly prefers. Pareto optimality has been termed the "single most important tool of normative economic analysis" (Moulin 2003). In case there is one object but multiple agents, there does not exist an assignment that is both Pareto optimal and envy-free. The reason is that a Pareto optimal assignment allocates the object to *some* agent in which case the other agents envy him.

In view of the importance of the two-agent setting and the fact that there may not exist a Pareto optimal and envy-free assignment, Brams et al. (2014) presented an elegant algorithm called AL for the case of two agents that computes a maximal assignment that is envy-free as well as *locally Pareto optimal* (Pareto optimal for the set of allocated objects).[1] The algorithm has received attention in the literature (see e.g., Brams 2014; Bouveret et al. 2015; Dickerson et al. 2014; Procaccia and Wang 2014). The desirable aspect of AL is that it returns a locally Pareto optimal and a maximal envy-free assignment. By maximal, we mean that unallocated objects cannot be added to the agents' partial allocations without compromising envy-freeness. Brams et al. (2014) also claim that AL returns a complete envy-free assignment if there exists one complete envy-free assignment. One possible limitation of the AL method is that it assumes that agents have strict preferences over objects. We present a generalization of the AL method in which agents may express indifferences among objects.

Indifferences in preferences are not only a natural relaxation but are also a practical reality in many cases. For example, if there are multiple copies of the same object with

[1] The notion of envy-freeness that they use is equivalent to SD (stochastic dominance) envy-freeness (Aziz et al. 2014) and necessary envy-freeness (Bouveret et al. 2010).

the same characteristics, then an agent is invariably indifferent among all such copies. Indifferences can lead to various challenges. The complexity of solution concepts in the presence of indifferences can be considerably more than in the case of strict preferences. A famous example is that of roommate markets for which the problem of finding a stable matching is polynomial time solvable for strict preferences but NP-complete for weak orders (Ronn 1990). Similarly, a number of fairness concepts are harder to compute when weak orders are allowed (Aziz et al. 2014). In view of this, effort has been taken to generalize algorithms and rules for the case of indifferences in voting (see e.g., Aziz et al. 2013a; Cullinan et al. 2014), housing markets (see e.g., Aziz and de Keijzer 2012; Saban and Sethuraman 2013), coalition formation (Aziz et al. 2013b), and various matching market models (Iwama and Miyazaki 2008; Manlove 2013; Scott 2005). The main contribution of this paper is a generalization of AL which we refer to as GAL for the case in which agents may express indifferences. The main result of the paper is as follows.

Theorem 1 *For two agents, GAL returns in time $O(m^2)$ a maximal envy-free and locally Pareto optimal assignment even if agents express weak preferences. If a complete envy-free assignment exists, GAL computes a complete envy-free assignment. Moreover, there exists no other assignment that Pareto dominates it and is envy-free and there exists no other assignment that allocates a superset of the allocated objects and is envy-free.*

Previously, Bouveret et al. (2010) and Aziz et al. (2014) presented $O(m^3)$ time algorithms to check whether a *complete* envy-free assignment for two agents exists or not. The algorithms require solving network flow or maximum matching problems. As a corollary of GAL, we obtain a simple $O(m^2)$ algorithm to check whether there exists a complete assignment that is EF.

A critical reader may ask whether GAL can be generalized to handle an arbitrary number of agents. We show that unless complexity classes P and NP coincide (Fortnow 2013), there exists *no* polynomial time algorithm for an arbitrary number of agents that satisfies the same properties as GAL.

2 Related work

Computation of fair discrete assignments has been intensely studied in the last decade. In many of the papers considered, agents express cardinal utilities for the objects and the goal is to compute fair assignments (see e.g., Lipton et al. 2004; Procaccia and Wang 2014). We consider the setting in which agents only express ordinal preferences over objects (Aziz et al. 2014; Bouveret et al. 2010; Brams and Kaplan 2004; Brams et al. 2003; Brams and Fishburn 2000; Pruhs and Woeginger 2012) which are less demanding to elicit.

When agents express preferences over objects and we need to reason about preferences over allocations, there are different ways one can define envy-freeness such as possible envy-freeness and weak SD envy-freeness (Aziz et al. 2014). In this paper, we will use the strongest known reasonable notion of envy-freeness. The notion is equivalent to ‘not envy-possible’ (Brams et al. 2003), necessary envy-freeness (Bouveret et al. 2010), SD envy-freeness (Aziz et al. 2014), and EF (envy-freeness) notion used by Brams et al. (2014). We will refer to the notion simply as EF just like Brams et al. (2014) do. Aziz et al. (2014) and Bouveret et al. (2010) presented $O(m^3)$ algorithms to check whether there exists a complete EF assignment. We show that there exists a simple $O(m^2)$ algorithm for the problem even if agents express weak preferences.

There are other papers (Aziz 2014; Chevaleyre et al. 2006; de Keijzer et al. 2009) in fair division in which agents explicitly express ordinal preferences over sets of objects rather than simply expressing preferences over objects (Aziz 2014; Chevaleyre et al. 2006; de Keijzer et al. 2009). For these more expressive models, the computational complexity of computing fair assignments is either even higher (Chevaleyre et al. 2006; de Keijzer et al. 2009) or representing preferences requires exponential space (Aziz 2014; Brams et al. 2012). In this paper, we restrict agents to simply express ordinal preferences over objects.

3 Preliminaries

An assignment problem is a triple $(N, O, \succsim)$ such that $N = \{1, \ldots, n\}$ is a set of agents, $O = \{o_1, \ldots, o_m\}$ is a set of objects, and the preference profile $\succsim = (\succsim_1, \ldots, \succsim_n)$ specifies for each agent i its preference $\succsim_i$ over O. Agents may be indifferent among objects. We will denote by $\succ_i$ the strict part and by $\sim_i$ the indifference part of the relation $\succsim_i$. We denote $\succsim_i$: $E_i^1, \ldots, E_i^{k_i}$ for each agent i with equivalence classes in decreasing order of preferences. Thus, each set E_i^j is a maximal equivalence class of objects among which agent i is indifferent, and k_i is the number of equivalence classes of agent i. If an equivalence class is a singleton $\{o\}$, we list the object o in the list without the curly brackets. A preference profile consists of *dichotomous preferences* if each agent has at most two equivalence classes. A preference profile consists of *strict preferences* if each agent has strict preferences over the objects. For a subset of objects $O' \subseteq O$, we will denote $\{o \in O' : o \succsim_i o' \in O'\}$ by $\max_{\succsim_i}(O')$ and $\{o \in O' : o \precsim_i o' \in O'\}$ by $\min_{\succsim_i}(O')$.

Definition 1 (*Assignment*) An assignment $p = (p(1), \ldots, p(n))$ specifies the *allocation* of objects $p(i)$ to each $i \in N$ such that $p(i) \subseteq O$ and $p(i) \cap p(j) = \emptyset$ for all $i \neq j$. An assignment is *complete* if

$$\bigcup_{i \in N} p(i) = O.$$

We define the stochastic dominance (SD) relation which extends preferences over objects to preferences over sets of objects. An agent *SD prefers* one allocation over another if for each object o, the former allocation gives the agent as many objects at least as preferred as o as the latter allocation.

Definition 2 [(*SD* (*stochastic dominance*)] Given two 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 for each $o \in O$,

$$\left|\{o' : o' \succsim_i o\} \cap p(i)\right| \geq \left|\{o' : o' \succsim_i o\} \cap q(i)\right|.$$

We say that agent i *strictly SD prefers* $p(i)$ to $q(i)$ denoted by $p(i) \succ_i^{SD} q(i)$ if $p(i) \succsim_i^{SD} q(i)$ and $\neg[q(i) \succsim_i^{SD} p(i)]$.

Although each agent i expresses ordinal preferences over objects, he could have a private cardinal utility u_i consistent with $\succsim_i$: $u_i(o) \geq u_i(o')$ if and only if $o \succsim_i o'$.

Definition 3 (*SD-efficiency*) An assignment p is *SD-efficient* if there exists no other assignment q such that $q(i) \succsim_i^{SD} p(i)$ for all $i \in N$ and $q(i) \succ_i^{SD} p(i)$ for some $i \in N$.

SD-efficiency is equivalent to Pareto optimality for discrete assignments as defined by Brams et al. (2014). Hence, we will refer to SD-efficiency as Pareto optimality and SD-domination as Pareto domination.

Definition 4 (*Locally Pareto optimal*) An assignment p is *LPO (locally Pareto optimal)* if there exists no other assignment q such that $\bigcup_{i\in N} p(i) = \bigcup_{i\in N} q(i)$ and $q(i) \succsim_i^{SD} p(i)$ for all $i \in N$ and $q(i) \succ_i^{SD} p(i)$ for some $i \in N$.

Definition 5 (*SD envy-freeness*) An assignment p satisfies *SD envy-freeness* or is *SD envy-free* if each agent SD prefers its allocation to that of any other agent:

$$p(i) \succsim_i^{SD} p(j) \quad \text{for all } i, j \in N.$$

From the definition, it is easy to see that a necessary condition for SD envy-freeness is that each agent gets the same number of objects.

Brams et al. (2014) defined EF as follows. We use the same definition as the one by Brams et al. (2014) but write it for any number of agents. Furthermore, we use weakly prefers rather than strictly prefers since we are considering weak preferences.

Definition 6 [*EF* (*envy-freeness*)] An allocation p is *EF (envy-free)* if for all $i, j \in N$, $|p(i)| = |p(j)|$ and there exists an injection $f_{ij} : p(i) \to p(j)$ and an injection $f_{ji} : p(j) \to p(i)$ such that for each object $o \in p(i)$, i (weakly) prefers o to $f_{ij}(o)$ and for each object $o \in p(j)$, j (weakly) prefers o to $f_{ji}(o)$.

Therefore, using a similar argument as (Lemma 1, Brams et al. 2014), we can show that EF is equivalent to SD envy-freeness. We detail the argument for the sake of completeness and to formally extend Lemma 1(Brams et al. 2014) to the case of indifferences.

Lemma 1 *EF is equivalent to SD envy-freeness.*

Proof We first show that EF implies SD envy-freeness. Suppose p satisfies EF and take any object $o \in O$. Suppose that there is an object $o' \in p(i)$ such that $f_{ij}(o') \succsim_i o$. By the definition of f_{ij}, we know that $o' \succsim_i f_{ij}(o')$. Since $f_{ij}(o') \succsim_i o$, we get that $o' \succsim_i f_{ij}(o') \succsim_i o$. Hence,

$$\left|\{o' : o' \succsim_i o\} \cap p(i)\right| \geq \left|\{o' : o' \succsim_i o\} \cap p(j)\right|.$$

We now show that SD envy-freeness implies EF. Suppose that assignment p does not satisfy EF with the EF condition violated for agent i. Consider a bipartite graph $G = (p(j) \cup p(i), E)$ where $\{o, o'\} \in E$ if $o \in p(j)$, $o' \in p(i)$, and $o' \succsim_i o$. Since p does not satisfy EF for i, G does not admit a perfect matching. By Hall's theorem, there exists set $O' \subseteq p(j)$ such that $|N(O')| < |O'|$ where N is the neighborhood of O' in the graph G. Consider an object $o \in \min_{\succsim_i}(O')$. Since, $|N(O')| < |O'|$, this it implies that

$$\left|\{o' : o' \succsim_i o\} \cap p(i)\right| < \left|\{o' : o' \succsim_i o\} \cap p(j)\right|.$$

But then p does not satisfy SD envy-freeness. □

Lemma 2 *It can be checked in $O(n^2m)$ time whether a given assignment is EF or not.*

Proof We show that it can be checked in $O(m)$ time whether a given assignment for a constant number agents is SD envy-free or not. We first show that an SD comparison between any two allocations can be made in $O(m)$ time. Let us say that we want to check whether $p(i) \succsim_i^{SD} p(j)$ where j is some agent other than i. Without loss of generality, assume that i's preferences are a coarsening of linear order $o_1, \ldots, o_m$.

- We construct in $O(m)$ a vector $x(p(i)) = (x_1, \ldots, x_m)$ where $x_i = 1$ if $o_i \in p(i)$ and $x_i = 0$ otherwise. Using $x(p(i))$, we construct in $O(k_i)$ time a vector $s'(p(i)) = (s'_1, \ldots, s'_{k_i})$ where $s'_j = |\{E_i^j\} \cap p(i)|$. Using $s'(p(i))$, we construct in $O(k_i)$ time a vector $s(p(i)) = (s_1, \ldots, s_{k_i})$ where $s_j = \sum_{\ell=1}^{j} s'_j$.
- In a similar way, we construct in $O(m)$ a vector $y(p(j)) = (y_1, \ldots, y_m)$ where $y_i = 1$ if $o_i \in p(j)$ and $x_i = 0$ otherwise. Using $y(p(i))$, we construct in $O(k_i)$ time a vector $t'(p(j)) = (t'_1, \ldots, t'_{k_i})$ where $t_j = |\{E_i^j\} \cap p(j)|$. Using $t'(p(j))$, we construct in $O(k_i)$ time a vector $t(p(i)) = (t_1, \ldots, t_{k_i})$ where $t_j = \sum_{\ell=1}^{j} t'_\ell$.

Now $p(i) \succsim_i^{SD} p(j)$ iff $s_j \geq t_j$ for all $j \in \{1, \ldots, k_i\}$. This again takes time $O(k_i)$. Hence, an SD comparison between allocations takes time $O(m) + 4O(k_i) = O(m)$.

To test EF, we need to make $n(n-1)$ comparisons. Hence, testing EF of an assignment takes time $O(n^2m)$. □

In the paper, we will assume that $n = 2$, i.e., there are two agents. Hence for $n = 2$, EF can be tested in time $O(m)$. If we refer to some agent as $i \in \{1, 2\}$, then we will refer to the other agent as $-i$. Even for more than two agents, we may refer to $-i$ as some agents other than $i \in N$.

We define *maximal envy-freeness*.

Definition 7 (*Maximally envy-free assignment*) We say that a partial assignment p is *maximally envy-free* if it is envy-free and there exists no assignment q such that $q(i) \supseteq p(i)$ for all $i \in N$, $q(i) \supset p(i)$ for some $i \in N$, and q is envy-free.

Finally, we define *strong maximal envy-freeness*.

Definition 8 (*Strong maximally envy-free assignment*) We say that a partial assignment p is *strong maximally envy-free* if it is envy-free and there exists no assignment q such that $\bigcup_{i \in N} q(i) \supset \bigcup_{i \in N} p(i)$ and q is envy-free.

We note that strong maximal envy-freeness implies maximal envy-freeness.

4 GAL: Generalized AL

Before we delve into GAL, we first informally describe a simplified version of AL that still satisfies the properties of AL as described by Brams et al. (2014). Agents have strict preferences and, in each round, they pick one object each. The algorithm repeats the following until all objects have been allocated to agent 1, 2, or the contested pile C. We will refer to an object as *unallocated* if it has not been allocated to 1 or 2 or placed in C. If the most preferred unallocated object of the agents is not the same, each agent picks its most preferred object. Otherwise, if the most preferred unallocated object o coincides, then we check whether we can give it to agent 1. If o is given to agent 1 and the next most preferred unallocated object is given to agent 2 and the partial assignment satisfies EF, then we allow such an allocation in the round. If not, we check in the same way whether we can give it to agent 2.[2] If o cannot be given to either of the two, we put it in C.

The general idea of GAL is as follows. Since the preferences of the two agents are weak orders, we first construct unique linear orders called *priority orders* based on the preferences. Although the comparisons to check the feasibility of EF assignments are still done with respect to the original preferences, the constructed linear orders help to identify which unique object should each agent try to get first. The priority orders are refinements of the preferences where, if an agent is indifferent between two objects, it has higher priority for the object less preferred by the other agent. If both agents are indifferent among two objects, then agent 1 has higher priority for the object with the lower index and agent 2 has higher priority for the object with the higher index. After suitably constructing the linear orders, $>_1$ and $>_2$, agents try to take the highest priority object. If agents have a different highest priority object, they take their highest priority objects. Otherwise, there is a conflict so we must try to give one of the agents the highest priority object and give the other agent the second highest priority object according to the priority list if it does not violate EF.[3] If this cannot be done, we send the contested object to C, the so-called *contested pile*. A key idea behind GAL is that if an object o^* is sent to the contested pile, then it cannot be the case that o^* along with some subsequent less preferred objects are allocated to agents and EF is not violated. The algorithm is formally defined as Algorithm 1. Note that there is an asymmetry in the algorithm in that agent one is considered first to get object o^* in Step 18. One can consider any of the two agents first or even toss a coin to select one agent. The properties of the algorithm are not affected.

[2] This feasibility check is phrased in a different way in the original description of AL but is equivalent to checking for EF.

[3] The view of EF as being defined with respect to the SD relation makes it easy to argue for a maximal EF assignment.

Algorithm 1 GAL—algorithm for envy-free assignment of indivisible objects to two agents

Input: $((\succsim_1, \succsim_2), O)$
Output: EF assignment p

Construct the linear order $>_1$ for agent 1: for all $i, j \in \{1, \ldots, m\}$, $o_i >_1 o_j$ if $o_i \succ_1 o_j$; $o_j >_1 o_i$ if $o_i \succ_2 o_j$ and $o_i \sim_1 o_j$; $o_i >_1 o_j$ if $o_i \sim_2 o_j$ and $o_i \sim_1 o_j$ and $i < j$.
Construct the linear order $>_2$ for agent 2: for all $i, j \in \{1, \ldots, m\}$, $o_i >_2 o_j$ if $o_i \succ_2 o_j$; $o_j >_2 o_i$ if $o_i \succ_1 o_j$ and $o_i \sim_2 o_j$; $o_j >_2 o_i$ if $o_i \sim_2 o_j$ and $o_i \sim_1 o_j$ and $i < j$.
$O' \longleftarrow O$
$p(1) \longleftarrow \emptyset$; $p(2) \longleftarrow \emptyset$
$C \longleftarrow \emptyset$
round number $t \longleftarrow 0$
while $O' \neq \emptyset$ **do**
  $t \longleftarrow t + 1$
  if $|O'| = 1$ **then**
    $C \longleftarrow C \cup O'$
  else if $\max_{>_1}(O') \neq \max_{>_2}(O')$ **then**
    $p(1) \longleftarrow p(1) \cup \max_{>_1}(O')$
    $p(2) \longleftarrow p(2) \cup \max_{>_2}(O')$
    $O' \longleftarrow O' \setminus \{\max_{>_1}(O'), \max_{>_2}(O')\}$
  else if $\max_{>_1}(O') = \max_{>_2}(O')$ **then**
    $o^* \longleftarrow \max_{>_1}(O')$ (or $\max_{>_2}(O')$)
    $O' \longleftarrow O' \setminus \{o^*\}$
    if $(p(1) \cup \{o^*\}, p(2) \cup \{\max_{>_2}(O')\})$ is EF w.r.t $\succsim$ **then**
      $p(1) \longleftarrow p(1) \cup \{o^*\}$
      $p(2) \longleftarrow p(2) \cup \{\max_{>_2}(O')\}$
      $O' \longleftarrow O' \setminus \{\max_{>_2}(O')\}$
    else if $(p(1) \cup \{\max_{>_1}(O')\}, p(2) \cup \{o^*\})$ is EF w.r.t $\succsim$ **then**
      $p(2) \longleftarrow p(2) \cup \{o^*\}$
      $p(1) \longleftarrow p(1) \cup \{\max_{>_1}(O')\}$
      $O' \longleftarrow O' \setminus \{\{\max_{>_1}(O')\}$
    else
      $C \longleftarrow C \cup \{o^*\}$
    end if
  end if
end while
return $(p(1), p(2))$

First observe that for strict preferences, GAL is equivalent to the simplified AL method. The reason is that for strict preferences, there exists a unique priority order irrespective of any lexicographical tie-breaking order. We present a couple of examples to illustrate how GAL works. The contested pile is empty in one example and non-empty in another.

Example 1

$$\begin{aligned} \succsim_1&: \{o_1, o_2, o_3\}, \{o_4, o_5, o_6\} \\ \succsim_2&: \{o_2, o_3, o_4\}, \{o_6\}, \{o_1, o_5\} \\ >_1&: o_1, o_2, o_3, o_5, o_6, o_4 \\ >_2&: o_4, o_3, o_2, o_6, o_5, o_1 \end{aligned}$$

(i) Round 1: $p(1) = \{o_1\}$, $p(2) = \{o_4\}$, $C = \emptyset$;

(ii) Round 2: $p(1) = \{o_1, o_2\}$, $p(2) = \{o_4, o_3\}$, $C = \emptyset$;
(iii) Round 3: $p(1) = \{o_1, o_2, o_5\}$, $p(2) = \{o_4, o_3, o_6\}$, $C = \emptyset$.

Example 2

$$\begin{aligned}
\succsim_1: &\quad \{o_7\}, \{o_1, o_2, o_3\}, \{o_4, o_5, o_6\} \\
\succsim_2: &\quad \{o_7\}, \{o_1\}, \{o_3\}, \{o_4, o_5\}, \{o_2, o_6\} \\
>_1: &\quad o_7, o_2, o_3, o_1, o_6, o_4, o_5 \\
>_2: &\quad o_7, o_1, o_3, o_5, o_4, o_6, o_2
\end{aligned}$$

(i) Round 1: $p(1) = \emptyset$, $p(2) = \emptyset$, $C = \{o_7\}$;
(ii) Round 2: $p(1) = \{o_2\}$, $p(2) = \{o_1\}$, $C = \{o_7\}$;
(iii) Round 3: $p(1) = \{o_2, o_3\}$, $p(2) = \{o_1, o_5\}$, $C = \{o_7\}$;
(iv) Round 4: $p(1) = \{o_2, o_3, o_6\}$, $p(2) = \{o_1, o_5, o_4\}$, $C = \{o_7\}$.

Proposition 1 *GAL runs in $O(m^2)$ time and is deterministic.*

Proof The priority orders of the agents can be computed in time $O(m^2)$. The weak preferences $\succsim_i$ of each agent are viewed as an adjacency matrix M_i in which $M_i[o, o'] = 1$ if $o \succ_i o'$, $M_i[o, o'] = 0$ if $o \sim_i o'$, and $M_i[o, o'] = -1$ if $o \prec_i o'$. The matrix can be computed in time $O(m^2)$. For each pair $o, o' \in O$ in the priority list of agent i, the relative positions can be decided by examining the two preference comparisons $\succsim_i$ and $\succsim_i'$ between o and o'. Hence, the priority order of an agent can be computed in $O(m^2)$.

In each round, either one object each is allocated to the agents or one contested object is sent to C. If each agent has a different highest priority unallocated object, then the allocation takes constant time. Otherwise, the agents have the same highest priority contested object o^*. In this case, we need to make at most two checks for whether there exists an EF partial assignment that allocates o^* to one of the agents. In either of these checks, we simply need to verify whether the given partial assignment is EF or not which takes time $O(m)$ according to Lemma 2. Thus, GAL takes time $O(m^2)$. □

We present a simple argument why GAL returns a maximal EF assignment. Previously, Brams et al. (2014) had proved that AL returns a maximal EF assignment when preferences are strict. Later, we will prove that GAL even satisfies strong maximal envy-freeness.

Proposition 2 *GAL returns a maximal EF assignment.*

Proof The GAL outcome is EF. This follows from the way the partial assignments are constructed so that EF is maintained. Consider O' the set of unallocated objects. If $\max_{>_1}(O') = \max_{>_2}(O')$, then the partial assignment is only modified after checking that the modification still satisfies EF. If $\max_{>_1}(O') \neq \max_{>_1}(O')$, then each agent is given a most preferred unallocated object from O'. Since the partial allocation p is EF, and for each $o \in p(i)$, $o \succsim_i \max_{>_i}(O') \succsim_i \max_{>_{-i}}(O')$, it follows that the allocation which gives $p(i) \cup \{\max_{>_i}(O')\}$ to each $i \in \{1, 2\}$ is EF.

We now show that the outcome is a maximal EF assignment. Assume for contradiction that GAL's outcome p is not maximal EF. This means that for some object $o \in C$ there exists an assignment q that matches the objects matched by p as well as o and possibly other objects. Consider the first such object o that is to be placed in the contested pile C and consider the stage in Algorithm 1 where o was sent to C. If o was given to agent $-i$, then agent i was given the next highest priority object o' according to $>_i$ which still leads to infeasibility of EF. Clearly, $o \succ_i o'$ or else the partial assignment p^t at the stage would not fail EF. For every other unallocated object o'' in O' (that is not in the contested pile) at that stage, it holds that $o' \succsim_i o''$. Hence, no object o'' can be given to agent i while o is given to $-i$ so that p^t is still EF. Hence q cannot be EF, a contradiction. □

Next, we strengthen Proposition 2 by showing that GAL satisfies strong maximal envy-freeness. For the proposition, we require the following lemma. It says that a partial assignment is EF if for each allocated object o, an agent gets at least half of the allocated objects that are at least as preferred as o.

Lemma 3 *For the case of two agents, any partial assignment p is EF iff for each $o \in p(i) \cup p(-i)$,*

$$|p(i) \cap \{o' : o' \succsim_i o\}| \geq |\{o' \in p(i) \cup p(-i) : o' \succsim_i o\}|/2.$$

Proof For two agents, a partial assignment is EF if and only if it is SD-proportional (Theorem 4(ii), Aziz et al. 2015). Assignment p is SD-proportional if for each $i \in N$ and $o \in \cup_{i \in N} p(i)$,

$$|p(i) \cap \{o' : o' \succsim_i o\}| \geq |\{o' \in \cup_{i \in N} p(i) : o' \succsim_i o\}|/n.$$

For $n = 2$, we get the statement in the lemma. □

Proposition 3 *GAL returns a strong maximal EF assignment.*

Proof Let p be the partial assignment returned by GAL. Let the objects that are not allocated by p be $o'_1, o'_2, \ldots$ in the order in which they are placed in C during the running of GAL. We denote this order by π. Assume for contradiction that there exists an assignment q such that $\bigcup_{i \in N} q(i) \supset \bigcup_{i \in N} p(i)$ and q is envy-free. Without loss of generality, let q be such an assignment that allocates the earliest object o in the sequence π. This implies that there exists no EF assignment r such that $(p(i) \cup p(-i) \cup \{o'\}) \subset (r(i) \cup r(-i))$ where o' is an object that comes earlier than o in the sequence π.

When o is placed in the contested pile, let the partial EF allocation be p^t. We note that $o' \succsim_i o$ for all $o' \in p^t(i)$ and $o' \succsim_{-i} o$ for all $o' \in p^t(-i)$. If there are no objects left to be considered for allocation, then this is a contradiction since q allocates objects allocated by p as well as o without violating EF. Let us assume that there are still some objects to be considered for allocation. Let the next priority available unallocated objects of i and $-i$ be o^i and o^{-i}, respectively, where it could be possible that $o^i = o^{-i}$. Since o is placed in the contested pile, it follows that the assignment in which $p^t(i) \cup \{o\}$ is allocated to i and $p^t(-i) \cup \{o^{-i}\}$ is allocated to $-i$ not EF. Similarly, the assignment in which $p^t(i) \cup \{o^i\}$ is allocated to i and $p^t(-i) \cup \{o\}$ is allocated to $-i$ is not EF. This implies that $o \succ_i o^i$ and $o \succ_{-i} o^{-i}$. This implies that o is strictly more preferred by i and $-i$ than all objects that are neither currently in C nor in $p^t(i) \cup p^t(-i)$. We denote the current C as C^t. Note that $(q(i) \cup q(-i)) \cap C^t = \emptyset$ because o is the first object in sequence π that is allocated to an agent by q but not by p.

Now consider the objects in $p^t(i)$. All objects in $p^t(i)$ have a better rank than o for agent i. Secondly, in agent i's priority list, if an object that is weakly preferred by agent i to o is not allocated to i, it is allocated to agent $-i$. Now agent i's allocation $p^t(i)$ is such that if o is given agent $-i$ and o^i to agent i, the assignment is not EF. By Lemma 3, this means that

$$\begin{aligned}&|(p^t(i) \cup \{o^i\}) \cap \{o' \in p^t(i) \cup p^t(-i) \cup \{o, o^i\} : o' \succsim_i o\}| \\ &\quad < |\{o' \in p^t(i) \cup p^t(-i) \cup \{o, o^i\} : o' \succsim_i o\}|/2.\end{aligned}$$

Since $o \succ_i o^i$, the statement above is equivalent to

$$\begin{aligned}&|p^t(i) \cap \{o' \in p^t(i) \cup p^t(-i) \cup \{o\} : o' \succsim_i o\}| \\ &\quad < |\{o' \in p^t(i) \cup p(-i) \cup \{o\} : o' \succsim_i o\}|/2.\end{aligned}$$

Since $-i$ cannot get o without the partial assignment p^t violating envy-freeness, we get the following condition for agent $-i$ by symmetry.

$$\begin{aligned}&|p^t(-i) \cap \{o' \in p^t(i) \cup p^t(-i) \cup \{o\} : o' \succsim_{-i} o\}| \\ &\quad < |\{o' \in p^t(i) \cup p^t(-i) \cup \{o\} : o' \succsim_{-i} o\}|/2\end{aligned}$$

Since o is strictly more preferred by i and $-i$ over all the other unallocated objects, it follows from the above inequalities that

$$\{o' \in O \setminus C^t : o' \succsim_i o\} = p^t(i) \cup p^t(-i) \cup \{o\}$$

and

$$\{o' \in O \setminus C^t : o' \succsim_{-i} o\} = p^t(i) \cup p^t(-i) \cup \{o\}.$$

Since $(q(i) \cup q(-i)) \cap C^t = \emptyset$, it follows that

$$\{o' \in (q(i) \cup q(-i)) : o' \succsim_i o\} = p^t(i) \cup p^t(-i) \cup \{o\}$$

and

$$\{o' \in (q(i) \cup q(-i)) : o' \succsim_{-i} o\} = p^t(i) \cup p^t(-i) \cup \{o\}.$$

Hence, $\{o' \in (q(i) \cup q(-i)) : o' \succsim_i o\} = \{o' \in (q(i) \cup q(-i)) : o' \succsim_{-i} o\}$. Note that $|\{o' \in (q(i) \cup q(-i)) : o' \succsim_i o\}| = |\{o' \in (q(i) \cup q(-i)) : o' \succsim_{-i} o\}|$ is odd since $|p^t(i) \cup p^t(-i)|$ is even. Now, if

$$|q(i) \cap \{o' \in (q(i) \cup q(-i)) : o' \succsim_i o\}| \geq |\{o' \in q(i) \cup q(-i) : o' \succsim_i o\}|/2,$$

it follows that

$$|q(-i) \cap \{o' \in (q(i) \cup q(-i)) : o' \succsim_{-i} o\}| < |\{o' \in q(i) \cup q(-i) : o' \succsim_{-i} o\}|/2.$$

By Lemma 3, this means that EF is violated for agent $-i$. Hence, we have shown that p is a strong maximal EF assignment. □

The argument for Proposition 3 can be used to show that GAL returns a strong maximal EF assignment p that allocates the *maximum* number of objects.

Proposition 4 *GAL returns an EF assignment p that allocates the maximum number of objects.*

Proof We know from Proposition 3 that GAL returns a strong maximal EF assignment. Let us denote this assignment by p. Assume for contradiction that there exists another strong maximal EF assignment q that allocates strictly more objects. But this means that q allocates some objects to the agents that p does not allocate to the agents. Let o be the first such object that is placed in the contested pile during the running of GAL that returns p. By the argument in the proof of Proposition 3, there exists no EF assignment that allocates the objects allocated by GAL up till that point as well as o and possibly some other objects. □

Proposition 3 implies that if there exists a complete EF assignment, then GAL returns a complete EF assignment. [4]

Proposition 5 *If there exists a complete EF assignment, then GAL returns a complete EF assignment.*

Proof Since GAL returns a strong maximal EF assignment (Proposition 3), it will return a complete EF assignment if such an assignment exists. □

Next, we show that the GAL outcome is LPO. Unlike Brams et al. (2014), we cannot use the characterization of Brams and King (2005) that if agents have *strict* preferences, any assignment as a result of sequential allocation is Pareto optimal. Hence, we need a lemma.

Let $(N, O, \succsim)$ be an assignment problem and p be a discrete assignment. We will create an auxiliary assignment problem and assignment where each agent is allocated exactly one object (see e.g., Aziz et al. 2015). The *clones* of an agent $i \in N$ are the agents in $N'_i = \{i_o : o \in O \text{ and } o \in p(i)\}$. The *cloned assignment problem* of $(N, O, \succsim)$ is $(N', O, \succsim')$ such that $N' = \bigcup_{i \in N} N'_i$. and for each $i_o \in N'$, $\succsim'_{i_o} = \succsim_i$. The *cloned assignment* of p is the discrete assignment p' in which $o \in p'(i_o)$ if $o \in p(i)$ and $o \notin p'(i_o)$ otherwise. The *cloned assignment* of p is the discrete

[4] The argument in Theorem 3 of Brams et al. (2014) only shows that for strict preferences, AL finds maximally EF assignment. It does not show that for strict preferences, AL efficiently computes a complete EF assignment if a complete EF assignment exists.

assignment p' in which each clone i_o gets at most one object o and it gets o only if $o \in p(i)$. A cloned assignment can easily be transformed back into the original assignment where each agent $i \in N$ is allocated all the objects assigned by p' to the clones of i.

Lemma 4 *An assignment for two agents is LPO iff there exist no objects o, o' such that o is allocated to i, o' is allocated to $-i$, $o' \succ_i o$ and $o \succsim_{-i} o'$.*

Proof By (Lemma 2, Aziz et al. 2015), an assignment is Pareto optimal if and only if its cloned assignment is Pareto optimal for the cloned assignment problem. Hence, we can restrict our attention to the cloned assignment and the cloned assignment setting. If the cloned assignment is Pareto optimal, the original assignment is Pareto optimal.

If the cloned assignment is *not* Pareto optimal, then there exists a 'trading cycle' in which each object points to its owner, each cloned agent in the cycle points to an object that is at least as preferred as its own object and at least one agent in the cycle points to a strictly more preferred object than the one it owns (Aziz and de Keijzer 2012).

Firstly, we claim that there exists no trading cycle consisting only of clones of one agent. Assume for contradiction that there exists a trading cycle consisting only of clones of the same agent. Then, there exists at least one object that is minimally preferred. The agent who points to this object also owns a minimally preferred object. Hence, each agent owns a minimally preferred object, and thus the cycle is not Pareto improving.

We now show that, if there exists a trading cycle, then there exists one which alternates between clones of the two agents. Consider any cycle which has the following path consisting of multiple clones of the same agent in succession:

$$o_{c_1} \to i_{c_1} \to o_{c_2} \to i_{c_2} \to \cdots \to o_{c_k} \to i_{c_k} \to o_{c_{k+1}} \to -i_{c_{k+1}}.$$

Since clones of each agent i have the identical preference, i_{c_1} also points directly to $o_{c_{k+1}}$. Hence, we know that there is also a path

$$o_{c_1} \to i_{c_1} \to o_{c_{k+1}} \to -i_{c_{k+1}}.$$

We now show that if there exists a trading cycle, if a clone of agent i points to a strictly more preferred object, then we can assume it is the clone of agent i who owns the least preferred object among all clones of i in the cycle. By the definition of trading cycle, at least one agent points to an object strictly more preferred than the object he owns. Assume that a clone of agent i gets a strictly more preferred object in the trading cycle. Let such a clone be i_j that points to o^*. Consider the clone i_1 of agent i who has the least preferred object among all clones of i in the cycle. We can assume without loss of generality that i_1 points to a strictly more preferred object than the one he owns. If this were not the case, then we know that i_j has a path to i_1 and i_1 also strictly prefers o^* over the object he owns. This means that there is a trading cycle in which i_1 points to a strictly more preferred object owned by a clone of $-i$. Hence, without loss of generality, let the agents in the trading cycle have the following sequence where i_1 points to and strictly prefers the object of $-i_1$ over his own object:

$$i_1, -i_1, i_2, -i_2, \ldots, i_k, -i_k, i_1.$$

We now show that if there exists a trading cycle which alternates between clones of the two agents, then there exists one with exactly one clone of each agent. Note that $-i_1$ at least weakly prefers the object of i_2 to his own object. We distinguish between two cases: (i) Clone i_2 is indifferent between his object and the object owned by i_1 (ii) Clone i_2 strictly prefers his object over the object owned by i_1. In case of (i), clone i_2 strictly prefers $-i_1$'s object over his own object. But this means that $-i_1$ weakly prefers i_2's object over his own object and i_2 strictly prefers $-i_1$'s object in which case we have already shown that there exist o, o' such that o is allocated to i, o' is allocated to $-i$, $o' \succ_i o$ and $o \succsim_{-i} o'$. Now let us consider case (ii). Since i_1 has the least preferred object among all clones of i, it points to any object that i_2 points to. Hence, $-i_1$ and i_2 can be removed from the cycle, and we still have a trading cycle in which i_1 now points to the object of $-i_2$. By repeating the argument inductively, we can find the clone $-i_k$ who points to the object of i_1 and i_1 strictly prefers $-i_k$'s object over his own object. Hence, we have shown that there exist o, o' such that o is allocated to i, o' is allocated to $-i$, $o' \succ_i o$ and $o \succsim_{-i} o'$. □

We use Lemma 4 to obtain the following proposition.

Proposition 6 *The GAL outcome is LPO.*

Proof Let us constrain ourselves to the set of objects $O' \subseteq O$ that are allocated to agents 1 and 2. Now assume that the GAL outcome is not LPO. Then, the assignment with respect to O' is not PO. By Lemma 4, there exists $i \in \{1, 2\}$ such that i gets o in some round t, $o' \succ_i o$, where o' was allocated to $-i$ and $o \succsim_{-i} o'$. This means that o' was allocated to $-i$ in round $t' \leq t$. Now if $o \succ_{-i} o'$, then o would be a higher priority object for $-i$ so that it would not have gone for o' before o. Then, it must be that $o \sim_{-i} o'$. But, if $o \sim_{-i} o'$, then o would again be a higher priority object for $-i$ so that it would not have gone for o' before o, hence a contradiction. □

In Proposition 6, we showed that there exists no other (not necessarily EF) assignment that uses the same objects as the GAL outcome and is a Pareto improvement over the GAL outcome. Next, we show that there exists no other EF assignment that may use any objects and is a Pareto improvement over the GAL outcome.

Proposition 7 *GAL returns an assignment such that there exists no other assignment that Pareto dominates it and is EF.*

Proof Assume for contradiction that GAL's outcome p is SD-dominated by another EF assignment q such that $q(i) \succsim_i^{SD} p(i)$ for both i and $q(i) \succ_i^{SD} p(i)$ for at least one i. Since p is a GAL outcome and q is EF, q cannot allocate more objects to the agents than p. Let p^t be the partial EF assignment after round t of GAL. Let q^t be the partial assignment in which each agent i gets the t' most preferred objects among $q(i)$ where t' is the number of objects allocated to i in p^t. We now proceed in rounds where in each round we check the highest priority allocated object of each of the two agents that have not been checked. We check the partial assignments p^t and q^t in each round t to see whether $q^t(i) \succ_i^{SD} p^t(i)$. Let us assume that $q^t(i) \succ_i^{SD} p^t(i)$

and $q^t(-i) \succsim_{-i}^{SD} p^t(-i)$ for the smallest possible t . If both $q^t(i) \succ_i^{SD} p^t(i)$ and $q^t(-i) \succ_{-i}^{SD} p^t(-i)$, then it means that in q both get higher priority objects than p in that round. This is a contradiction as GAL would allocate these higher priority and more preferred objects to the agents. Now assume that $q^t(i) \succ_i^{SD} p^t(i)$ and $q^t(-i) \sim_{-i}^{SD} p^t(-i)$. This means that agent $-i$ gets an equally preferred object and the other agent i gets a higher priority object. But this is again a contradiction, because GAL would have allocated the more preferred object to i in that round. □

Note that for the case of two agents, Aziz et al. (2014) presented a polynomial time algorithm to check whether a complete SD envy-free assignment exists or not. To compute a maximal SD envy-free assignment, one can consider different subsets $O' \subset O$ and check whether a complete SD envy-free assignment exists or not for O'. However, this approach would require checking exponential number of subsets.

We have already shown that GAL satisfies the desirable properties of AL on a more general domain. Next, we show that under strict preferences GAL returns an assignment that is a possible outcome of AL. In this sense, GAL is a 'proper' generalization of AL.

Proposition 8 *For strict preferences, GAL returns an AL outcome.*

Proof For strict preferences, there exists a unique priority order irrespective of any lexicographical tie-breaking order. We show that under strict preferences, GAL and AL handle all the cases in an equivalent manner.

Let us compare the formal definition of AL (pp 133–134, Brams et al. 2014) with the pseudocode of GAL. In AL, in stage t, the direction "If one unallocated item remains, place it in CP and stop" is equivalent to Steps 9 and 10 of Algorithm 1.

In AL, in stage t, the direction "If no unallocated items remain, stop." is equivalent to the stopping condition in the while loop of Algorithm 1.

If both agents have different most preferred (equivalent to highest priority since the preferences are strict) unallocated objects, then both GAL and AL behave in the same manner and give the most preferred objects to the agents. For AL, this direction is specified in the last sentence of the stage t.1.

Finally, both algorithms have a check for when both agents have the same most preferred objects with this check being in stage t.2 in the specification of the AL method. In AL, the most preferred available contested object i is tentatively given to the one of the agents. In the specification of Algorithm 1, the most preferred available object is also tentatively given to one of the agents. Since, in Algorithm 1, this object is referred to as o^*, we will refer it as o^* for both algorithms. Let us say that the agent who gets it is agent $-i$. The other agent i is tentatively given the next most preferred object that is not yet allocated. In the description of AL, i could be given an even less preferred unallocated object but in at least one instantiation of AL, i is tentatively given the next most preferred object that is still available. According to AL, such a tentative assignment is *feasible* as long as the number of objects assigned to $-i$ including o^* or put in the contested pile ("unassigned") that i prefers to the next most preferred unallocated object is at most t. This means that for the tentative assignment p, $|\{o' \in p(i) : o' \succ_i o^*\}| \geq |\{o' \in p(i) \cup p(-i) : o' \succ_i o^*\}|/2$. Since agent i's allocation from the previous round consists of objects strictly preferred over

o^*, this means that i is not envious of $-i$ in p as long as i was not envious of $-i$ in the previous round. Thus, in both algorithms, the tentative assignment in which the contested object is given to agent $-i$ and the next most preferred unallocated object is given to agent i is made permanent if the modification does not cause envy. Hence, the feasibility check in the case of AL is equivalent to checking whether the tentative new assignment is EF. If the tentative assignment is not EF for o^* given to either of the two agents, then GAL puts o^* in the contested pile. Similarly, AL puts the object in the contested pile (stage $t.5$). □

5 Discussion

In this paper, we presented GAL that is a generalization of the AL method of Brams et al. (2014) for the fair allocation of indivisible objects among two agents. A crucial advantage of extending AL to GAL is for the case in which agents have identical preferences. If agents have strict and identical preferences, then AL assigns all the objects to the contested pile. However, if the preferences are really coarse, such as when all objects are equally preferred, then GAL assigns $\lfloor m/2 \rfloor$ to each agent.

GAL can also be used as an algorithm to solve previously studied problems within fair division:

Theorem 2 *There exists a $O(m^2)$ algorithm to check whether there exists a complete assignment that is EF.*

Proof By Proposition 5, if there exists a complete EF assignment, GAL returns such an assignment. □

Previous algorithms to solve this problem take time $O(m^3)$ and require solving network flow or maximum matching problems (Bouveret et al. 2010; Aziz et al. 2014).

GAL is specifically designed for the case of two agents. This raises the question whether GAL can be generalized to an algorithm that returns a strong maximal EF assignment for any number of agents.

Theorem 3 *Assume there exists an algorithm $\mathcal{A}$ that returns a strong maximal EF assignment. Then, $\mathcal{A}$ does not take polynomial time assuming $P \neq NP$.*

Proof Checking whether there exists a complete EF assignment is NP-complete for strict preferences (Proposition 7, Bouveret et al. 2010). Furthermore, checking whether there exists a complete EF assignment is NP-complete for dichotomous preferences (Theorem 10, Aziz et al. 2014).

If $\mathcal{A}$ is polynomial time, then it can be used to compute a strong maximal EF assignment in polynomial time. If the assignment is complete, we know that there exists a complete EF assignment. If the assignment is not complete, then we know that there does not exist a complete EF assignment. Hence, $\mathcal{A}$ can solve an NP-complete problem in polynomial time. □

GAL can also be seen as a discrete version of the *probabilistic serial (PS) algorithm* (Bogomolnaia and Moulin 2001; Katta and Sethuraman 2006). PS is a fractional assignment algorithm that allocates a maximal subset of objects (in fact it allocates all the objects) and still achieves SD-efficiency and SD envy-freeness. In the discrete domain, GAL is an algorithm that returns a maximal assignment that is both SD-efficient and SD envy-free.

In this paper, we assumed that all objects are acceptable. The case where some objects may be unacceptable to an agent can be handled. If an object is unacceptable to both agents, it can be discarded from the outset. If an object is only acceptable to one agent, it will only be given to that agent.

It will be interesting to apply the approach of maximal EF to weaker notions of fairness (Aziz et al. 2014; Bouveret et al. 2010). The approach of finding a maximally fair assignment also applies to any setting in which there may not exist a complete fair assignment. Finally, extending GAL to the case of constant number of agents is left as future work.

Acknowledgments The author thanks Steven Brams for sharing the paper on the AL method with him. He also appreciates Sajid Aziz, Steven Brams, Christian Klamler, and Chun Ye for their useful feedback and to the reviewers for their detailed comments. NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

References

Abraham, D.J., Cechlárová, K., Manlove, D., Mehlhorn, K.: Pareto optimality in house allocation problems. In: Proceedings of the 16th International Symposium on Algorithms and Computation (ISAAC). Lecture Notes in Computer Science (LNCS), vol. 3341, pp. 1163–1175 (2005)

Aziz, H.: A note on the undercut procedure. In: Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), pp. 1361–1362 (2014)

Aziz, H., de Keijzer, B.: Housing markets with indifferences: a tale of two mechanisms. In: Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI), pp. 1249–1255 (2012)

Aziz, H., Brandt, F., Brill, M.: On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In: Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), pp. 455–462, IFAAMAS (2013a)

Aziz, H., Brandt, F., Harrenstein, P.: Pareto optimality in coalition formation. Games Econ. Behav. **82**, 562–581 (2013b)

Aziz, H., Gaspers, S., Mackenzie, S., Walsh, T.: Fair assignment of indivisible objects under ordinal preferences. In: Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), pp. 1305–1312 (2014)

Aziz, H., Gaspers, S., Mackenzie, S., Walsh, T.: Fair assignment of indivisible objects under ordinal preferences. Artif. Intell. **227**, 71–92 (2015)

Bogomolnaia, A., Moulin, H.: A new solution to the random assignment problem. J. Econ. Theory **100**(2), 295–328 (2001)

Bouveret, S., Endriss, U., Lang, J.: Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In: Proceedings of the 19th European Conference on Artificial Intelligence (ECAI), pp. 387–392 (2010)

Bouveret, S., Chevaleyre, Y., Maudet, N.: Fair allocation of indivisible goods. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.) Handbook of Computational Social Choice, chapter 12. Cambridge University Press, Cambridge (2015)

Brams, S.J.: Dispute Over Divorce or Inheritance? Try This 'Envy-Free' Algorithm. Huffington Post, New York (2014a)

Brams, S.J.: Dividing the Indivisible. Plus Magazine, Cambridge (2014b)

Brams, S.J., Fishburn, P.C.: Fair division of indivisible items between two people with identical preferences: envy-freeness, pareto-optimality, and equity. Soc. Choice Welf. **17**, 247–267 (2000)

Brams, S.J., Kaplan, T.R.: Dividing the indivisible procedures for allocating cabinet ministries to political parties in a parliamentary system. J. Theor. Polit. **16**(2), 143–173 (2004)
Brams, S.J., King, D.L.: Efficient fair division: help the worst off or avoid envy? Ration. Soc. **17**(4), 387–421 (2005)
Brams, S.J., Edelman, P., Fishburn, P.C.: Fair division of indivisible items. Theory Decis. **55**, 147–180 (2003)
Brams, S.J., Kilgour, D.M., Klamler, C.: The undercut procedure: an algorithm for the envy-free division of indivisible items. Soc. Choice Welf. **39**, 615–631 (2012)
Brams, S.J., Kilgour, D.M., Klamler, C.: Two-person fair division of indivisible items: an efficient, envy-free algorithm. Not. AMS **61**(2), 130–141 (2014)
Chevaleyre, Y., Dunne, P.E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-Aguilar, J.A., Sousa, P.: Issues in multiagent resource allocation. Informatica **30**, 3–31 (2006)
Cullinan, J., Hsiao, S.H., Polett, D.: A Borda count for partially ordered ballots. Soc. Choice Welf. **42**(4), 913–926 (2014)
de Keijzer, B., Bouveret, S., Klos, T., Zhang, Y.: On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In: Proceedings of the 1st International Conference on Algorithmic Decision Theory, pp. 98–110 (2009)
Demko, S., Hill, T.P.: Equitable distribution of indivisible objects. Math. Soc. Sci. **16**, 145–158 (1988)
Dickerson, J.P., Goldman, J., Karp, J., Procaccia, A.D., Sandholm, T.: The computational rise and fall of fairness. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI), pp. 1405–1411, AAAI Press (2014)
Foley, D.: Resource allocation and the public sector. Yale Econ. Essays **7**, 45–98 (1967)
Fortnow, L.: The Golden Ticket: P, NP, and the Search for the Impossible. Princeton University Press, Princeton (2013)
Gärdenfors, P.: Assignment problem based on ordinal preferences. Manag. Sci. **20**, 331–340 (1973)
Iwama, K., Miyazaki, S.: A survey of the stable marriage problem and its variants. In: Proceedings of the International Conference on Informatics Education and Research for Knowledge-Circulating Society, pp. 131–136, IEEE Computer Society (2008)
Katta, A.-K., Sethuraman, J.: A solution to the random assignment problem on the full preference domain. J. Econ. Theory **131**(1), 231–250 (2006)
Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Proceedings of the 5th ACM Conference on Electronic Commerce (ACM-EC), pp. 125–131, ACM Press (2004)
Manlove, D.: Algorithmics of Matching Under Preferences. World Scientific Publishing Company, Singapore (2013)
Moulin, H.: Fair Division and Collective Welfare. The MIT Press, Cambridge (2003)
Procaccia, A.D., Wang, J.: Fair enough: guaranteeing approximate maximin shares. In: Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC), pp. 675–692, ACM Press (2014)
Pruhs, K., Woeginger, G.J.: Divorcing made easy. In: Proceedings of FUN, number 7288. Lecture Notes in Computer Science (LNCS), pp. 305–314 (2012)
Ronn, E.: NP-complete stable matching problems. J. Algorithms **11**(2), 285–304 (1990)
Saban, D., Sethuraman, J.: House allocation with indifferences: a generalization and a unified view. In: Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC), pp. 803–820, ACM Press (2013)
Scott, S.: A Study of Stable Marriage Problems with Ties. PhD thesis, University of Glassgow (2005)
Wilson, L.: Assignment using choice lists. Oper. Res. Q. **28**(3), 569–578 (1977)
Young, H.P.: Dividing the indivisible. Am. Behav. Sci. **38**, 904–920 (1995)