Randomised Room Assignment-Rent Division

Lachlan Dufton and **Kate Larson**
Cheriton School of Computer Science
University of Waterloo
Waterloo, Canada
{ltdufton, klarson}@cs.uwaterloo.ca

Abstract

The room assignment-rent division problem allocates a heterogeneous set of indivisible items (e.g. rooms in a house) along with a share of some divisible item (e.g. the rent for the house), such that all items and resources are allocated without surplus or deficit, and each agent receives exactly one indivisible item. It is desirable to have envy-free outcomes but this is not possible for deterministic, truthful mechanisms. In this work we present truthful, randomised mechanisms for this problem, along with new measures of envy appropriate for non-deterministic mechanisms.

1 Introduction

The room assignment-rent division problem (RA-RD) [Su, 1999] is a classic problem in multiagent resource allocation and fair division. Consider a group of friends who will rent a house together. They must decide both who gets which room, and what share of the rent each person will pay. Each friend will want to be allocated just one room and there should be no surplus or deficit when meeting the total rent. Each individual has his or her own preferences on which room is best, such as preferring the largest room, or the room with the best view. More generally, this can be seen as a problem of allocating a set of indivisible, heterogeneous items (i.e. the rooms) along with a share of a divisible resource (i.e. the rent), such that all items are allocated and each agent gets exactly one indivisible item. The resources can have positive or negative utility.

While the real-estate setting is intuitive, this model of resource allocation can be applied to problems in other areas. In a job or task allocation setting, the indivisible items are tasks with some negative utility, and the divisible resource is some payment to be distributed among the workers upon completion of the tasks. The workers or processors can be considered the indivisible resources, with agents submitting work and covering some cost of maintaining the equipment.

In these settings, we are interested in more than just a Pareto-efficient allocation, but also some notion of fairness. In this paper we focus on *envy* and *envy-freeness* as measures of fairness. An allocation assigns a bundle to each agent, where a bundle is a single item along with some share of the divisible resource. For a particular allocation of bundles to agents, an agent is *envious* if it views another agent's bundle as strictly better than its own. An envy-free mechanism provides an allocation where no agent is envious. Brams and Taylor [1996] discuss envy-freeness and other measures of fairness in fair division.

In the RA-RD problem, deterministic, envy-free mechanisms are vulnerable to manipulation by the participating agents. Since envy-free allocations in this setting are Pareto-efficient with balanced agent transfers, this is a consequence of the impossibility result of Green and Laffont [1979]. As such, previous work on this problem has focussed on procedures that have full information about agent preferences.

We use randomisation to create new mechanisms that achieve envy-freeness in strategy-proof mechanisms. Randomisation has been a powerful technique for overcoming impossibility results in past work on social choice problems. For example, in other item allocation settings [Faltings, 2005], k-self-selection [Alon *et al.*, 2010], and voting protocols [Procaccia, 2010]. We also examine appropriate measures of qualities such as envy-freeness in randomised mechanisms. As previous work on the room assignment-rent division problem focusses on deterministic mechanisms, existing measures are not entirely suitable. For envy-freeness, we look at the probability a mechanism returns an envy-free outcome. Additionally we use the expected number of envy-free agents to consider what happens over all possible outcomes. We provide bounds for these measures in truthful mechanisms.

1.1 The Model

The room assignment-rent division problem assigns a set of indivisible, heterogeneous items (e.g. rooms in a shared house), M, to a set of agents, N such that all agents receive exactly one item and $|N| = |M|$. There is also some total quantity T of a divisible resource (e.g. rent) to be completely divided among the agents. This allocation and division is performed simultaneously. Each agent $i \in N$ has a value for each item $j \in M$, denoted as $v_i(j)$ (or equivalently, $v_{i,j}$), with the unit of the divisible resource as the numeraire.

We do not assume complete knowledge of agent types, so an RA-RD mechanism receives a vector of reported agent values $\bar{V} =< \bar{v}_1, \dots \bar{v}_n >$ and produces an allocation function, $f : N \to M$, and a division $R \in \mathbb{R}^n$. A valid f must be bijective so every agent receives one item, every item is assigned to one agent. Let r_i denote the share of divisible resource agent i receives, where $\sum_{i \in N} r_i = T$. To simplify the notation, we let $v_i(f) = v_i(f(i))$. Agents have quasi-linear utilities, so an agent's utility for an allocation (f, R) is $u_i(f, R) = v_i(f) + r_i$.

In this work we use randomised mechanisms for the RA-RD problem. A deterministic mechanism takes a vector $\bar{V}$ of reported types and returns a single outcome, (f, R). A randomised mechanism instead uses a probability distribution over outcomes and returns a single outcome, (f, R), according to this distribution. Agents are risk neutral, so an agent's expected utility for a random distribution over outcomes is the probability-weighted sum of its utility of each outcome. A deterministic mechanism is truthful, or dominant strategy incentive compatible (DSIC), if no agent can increase its own utility by misreporting its type, regardless of other agents' actions. Similarly, a randomised mechanism is truthful *in expectation* if no agent can increase its own *expected* utility by misreporting, regardless of other agents' actions.

1.2 Related Work

For the room assignment-rent division problem there have been a number of previous solutions for finding envy-free solutions while assuming complete knowledge of agent types. Su [1999] proves the existence of envy-free outcomes for this setting, along with an interactive algorithm based on Sperner's lemma that uses simple queries to the agents. Abudlkadoiroğlu, Sönmez and Ünver [2004] developed an envy-free auction method for determining the allocation and prices of rooms with any number of agents that guarantees non-negative pricing. Haake, Raith and Su [2002] provided a more general procedure without the restrictions that the number of objects must equal the number of agents and each agent must receive exactly one object. For the room assignment-rent division problem, an envy-free solution relies on truthful preferences of the agents. Unfortunately, no deterministic mechanism exists that is both envy-free and non-manipulable.

Sun and Yang [2003] achieved a strategy-proof and envy-free mechanism for a similar allocation problem, but has different restrictions on the allocation of the divisible resource. Instead of dividing a single quantity of some resource, each indivisible item has its own "compensation limit". This model and proof was generalised by Andersson and Svensson [2008], and Andersson [2009] for greater flexibility on the indivisible objects, and a proof of coalitional strategy-proofness. However, the use of an item-based compensation limit instead of a single budget of divisible resource that must be entirely allocated mean that these mechanisms are incompatible for the room assignment-rent division model of this paper.

In this paper we use randomisation to achieve strategy-proof outcomes that are not possible in deterministic mechanisms. Moulin and Bogomolnaia [2001] and later Kojima [2009] examined a randomised mechanism for a similar allocation problem to RA-RD, but with the restriction that agents have the same ordinal ranking and where individual preferences are distinguished by a private "acceptance threshold". These randomised, strategy-proof mechanisms were shown to achieve efficient and envy-free outcomes. These papers also discuss methods of evaluating randomised allocation procedures. In a more general, but related item allocation setting, the Green-Laffont impossibility theorem [Green and Laffont, 1979] shows that for heterogeneous item allocation, no mechanism is Pareto-efficient, DSIC, and strong budget balanced. Strong budget balance requires that all agents' payments sum to exactly zero, while in RA-RD payments must sum to exactly T. Work by Faltings [2005] provided a randomised allocation technique that achieves incentive compatibility and budget balance at the expense of allocative efficiency. The quality of this randomised mechanism is assessed by the loss of efficiency in generated problems.

1.3 Deterministic RA-RD Mechanisms

As has been shown in previous work [Haake *et al.*, 2002; Su, 1999], no truthful, envy-free mechanism exists for the RA-RD problem. This follows by the Green-Laffont impossibility theorem [Green and Laffont, 1979], as an envy-free allocation is an efficient allocation [Alkan *et al.*, 1991], and the sum of payments must be budget balanced (if $T \neq 0$, each room can be given an initial charge of $\frac{T}{n}$ to bring the budget to zero). As the truthful mechanism cannot guarantee efficiency when ensuring the divisible resource is entirely allocated, the mechanism cannot provide an envy-free outcome for all inputs.

2 Envy-Freeness in Randomised Mechanisms

An envy-free, deterministic allocation mechanism produces an outcome where no agent prefers another agent's allocated bundle to its own. For the RA-RD problem, an outcome (f, R) is envy-free if :

$$v_i(f(i)) + r_i \geq v_i(f(j)) + r_j \text{ , } \forall i, j \in N$$

This measurement states whether or not a single outcome is envy-free. When examining randomised mechanisms, which can produce several outcomes for a single input, this does not provide an appropriate comparison of mechanisms. For this problem, it is beneficial to consider measures of envy-freeness designed for randomised mechanisms. In a randomised mechanism, agent envy can be measured before the randomisation process (i.e. on the agent's lottery of outcomes), or on the final outcome. A simple extension of measuring envy to a randomised mechanism is to compare each agent's lottery of allocations, prior to the mechanism performing its random selection.

Definition 1. *For* ex ante envy-freeness*, no agent strictly prefers another agent's lottery over final outcomes. Let K be the set of allocations, and p_k is the probability of choosing $k \in K$, which has associated allocation and payment functions (f^k, R^k). That is, for all agents $i \in N$:*

$$\sum_{k \in K} p_k(v_i(f^k(i)) + r_i^k) \geq \sum_{k \in K} p_k(v_i(f^k(j)) + r_j^k) \text{ , } \forall j \in N$$

Ex ante envy-freeness is trivial to achieve in truthful mechanisms – simply randomise over all possible allocations with equal probability and give each agent $\frac{T}{n}$ of the divisible resource. This gives the same lottery for each agent regardless of reported type but generally provides poor final outcomes, with most or all agents envious in all outcomes. Because of this, we propose looking at envy-freeness in the actual outcomes, after the mechanism has performed the random selection. One measure is to look at which of these final outcomes are envy-free in the deterministic sense, and the probability of the mechanism producing such an outcome in the worst case.

Definition 2. *An outcome is envy-free if no agent values another agent's bundle higher than its own. The* guaranteed probability of envy-freeness *(GPEF) is the minimum probability a mechanism will produce an envy-free outcome, for any set of agents.*

The previous example that was *ex ante* envy-free has a GPEF of zero. For some sets of agents, it will never produce an envy-free outcome. Consider two agents that both prefer indivisible item 1. As the divisible resource is split evenly, whichever agent is assigned item 2 will be envious. This measure only considers the very best outcomes, where all agents are envy-free, and all other outcomes are assessed as valueless. For our third measure, we examine the level of envy, as the number of envious agents, in each of the possible outcomes.

Definition 3. *An envy-free agent is one who does not value another agent's bundle higher than its own in a particular allocation. The* expected number of envy-free agents *is the probability-weighted sum of the number of envy-free agents in each outcome of the mechanism for a particular input.*

In the example of two agents preferring the same item, the basic mechanism gives 1 expected envy-free agent, as both allocations would have one agent envious and one envy-free.

3 Randomised RA-RD Mechanisms

We now examine these new measures of envy-freeness in the RA-RD problem on mechanisms that are truthful *in expectation*. A mechanism is truthful in expectation if, irrespective of the actions of other agents, an agent's expected utility can not be increased by misreporting its type.

Lemma 1. *An RA-RD mechanism is truthful in expectation if (but not only if) each agent's* expected *share of the divisible resource, and probability of being assigned to each indivisible item is constant (independent of reported types).*

Proof. Let $p_{i,j}$ be the probability agent i is assigned item j, and $\bar{r}_i = E(r_i)$ be agent i's expected share of the divisible resource. The expected utility of agent $i \in N$ is calculated as: $E(u_i) = \sum_{j \in M} p_{i,j} v_i(j) + \bar{r}_i$. As all $p_{i,j}$ and $\bar{r}_i$ are constant with respect to the agent's bid/reported type, the agent's expected utility is constant and cannot be increased by misreporting. □

Note that these are not the necessary conditions for a truthful RA-RD mechanism. We use these conditions to define a simple, truthful mechanism as a baseline for comparing other randomised mechanisms.

A simple randomised RA-RD mechanism. From previous work [Alkan *et al.*, 1991; Haake *et al.*, 2002], given full knowledge of agents' types, we can find an envy-free allocation and division, denoted (f^*, R^*). Our random mechanism first calculates the envy-free solution, then randomly selects an integer $x \in [0, n-1]$. Agent i is given the item and share allocated to agent $(i + x) \pmod n$ in the envy-free allocation. Thus, $f^x(i) = f^*((i + x) \pmod n)$. Each agent has a $\frac{1}{n}$ probability of being assigned any particular item. An agent's expected share of the divisible resource is

$$\bar{r}_i = \sum_{j \in n} \frac{1}{n} r_j^* = \frac{1}{n} \sum_{j \in n} r_j^* = \frac{T}{n}$$

This is constant for each agent, so by Lemma 1 the mechanism is truthful in expectation, allowing the mechanism to correctly calculate (f^*, R^*).

Whenever $x = 0$, the envy-free outcome is chosen, and this occurs with probability $\frac{1}{n}$. Apart from special cases, for all other values of x, all agents will be envious of their bundle from the envy-free outcome. Thus, for this mechanism the GPEF is $\frac{1}{n}$. When $x = 0$, there are n envy-free agents, while in the worst case, all other choices of x will have no envy-free agents. This gives a worst-case expected number of envy-free agents of $n \cdot \frac{1}{n} + 0 \cdot \frac{n-1}{n} = 1$. In this mechanism, all agents have the same lottery over items and expected payment, so it is *ex ante* envy-free.

3.1 Maximising Guaranteed Probability of Envy-Freeness

A truthful mechanism that guarantees 100% probability of envy-freeness would be optimal for the three definitions in Section 2. Unfortunately, this is not possible for RA-RD.

Theorem 1. *A truthful (in expectation) mechanism for the RA-RD problem with n agents has a guaranteed probability of envy-freeness of* at most $\frac{1}{n}$.

Proof. In our setting with an equal number of agents and items, an envy-free allocation is a Pareto-efficient allocation [Alkan *et al.*, 1991]. So, if a mechanism were capable of envy-freeness with probability $p > \frac{1}{n}$, it would also provide an efficient allocation with probability at least p.

To get an efficient allocation with probability more than $\frac{1}{n}$, then all agents must be able to change their probabilities of item allocation through their reported values. For any mechanism for this problem, an agent's expected utility, which must be maximised when reporting truthfully, can be decomposed into parts. The first is its expected utility from receiving items – a probability-weighted sum of the resources it can receive. An agent will always receive one item. Next, the agent's *expected* payment for any mechanism can be separated into two functions $\bar{g}_i(v) + h_i(v_{-i})$. Function $\bar{g}_i(v)$ depends on all agents' reported types and must be maximised when agent i reports truthfully (similar to the Groves payment in a Vickrey-Clarke-Groves (VCG) mechanism). There is some additional expected payment, $h_i(v_{-i})$, that doesn't depend on agent i's reported type.

Let $p_{i,j}(v)$ denote the probability agent i receives item j. If an agent receives each item with equal probability, then the agent will receive some constant expected utility from the allocation, regardless of its reported type. The minimum probability an agent can receive an item, $\min_{i,v} p_{i,j}(v)$, determines the fraction of outcomes that contribute to this constant utility. All items must be received with probability *at least*

$\min_{i,v} p_{i,j}(v)$. So with n items, let $p_i^0 = n \cdot \min_{i,v} p_{i,j}(v)$ be the fraction of outcomes where each agent receives each item with equal probability. Reported values do not affect expected utility from these allocations and the contribution to $\bar{g}_i(v)$ to ensure truthfulness is 0.

If $p_i^*(v)$ is the probability an agent receives its item in the efficient allocation, then it receives this item with increased probability of $(p_i^*(v) - \frac{p_i^0}{n})$ over the equal-probability allocations that are independent of bids. For truthfulness, $\bar{g}_i(v)$ must include $(p_i^*(v) - \frac{p_i^0}{n})g_i(v)$, where $g_i(v) = \sum_{j \neq i} v_j(f^*(j))$ is the Groves payment. This maximises the agents expected utility when it bids such that the true efficient allocation is chosen. Finally, if $p_i^0(v) + p_i^*(v) < 1$, there are other, non-efficient allocations for which agent i can change the probability, but the mechanism cannot counteract any gain from misreporting without $\bar{g}_i(v)$ directly including agent i's reported values, which will allow agent i to benefit by reducing its payment. This gives an agent's final expected payment of $(1-p_i^0)g_i(v)+h_i(v_{-i})$. The efficient allocation is possible with probability at most $p_i^*(v) = (1 - p_i^0)$.

All agents' expected payments must sum to T. If $p_i^0 < 1$, then dividing the h functions by constant factor $(1 - p_i^0)$ would give budget balanced Groves transfers. This contradicts the Green-Laffont impossibility theorem, so $p_i^0 = 1$.

With constant probability of being assigned each item, an agent cannot change its expected utility from the allocation by misreporting. This limits the probability of an efficient allocation to at most $\frac{1}{n}$ in the worst case (where there is only a single efficient allocation). Envy-free outcomes have efficient allocations so the best GPEF is $\frac{1}{n}$. □

This is a tight bound as demonstrated by the simple randomised RA-RD mechanism described above, with a GPEF of $\frac{1}{n}$. This places some limiting restrictions on what is possible with a strategy-proof mechanism for this problem. Envy-freeness at a low probability that asymptotically goes to zero means that most of the time the mechanism will produce a "bad" result. Considering only envy-free outcomes ignores what happens in the remainder of cases. In the mechanism described above, in the $(n-1)$ non-envy-free outcomes, every single agent will be envious. This motivates measuring the quality of each outcome with more detail than a yes/no test of "envy-free".

3.2 Maximising Expected Number of Envy-Free Agents

While having all agents envy-free is the ideal outcome, attempting to maximise the probability of such an outcome can come at the expense of the quality of non-envy-free outcomes. For truthful mechanisms, these non-envy-free outcomes are the most likely, so when comparing mechanisms they should not be ignored.

The above mechanism with a GPEF of $\frac{1}{n}$ has expected number of envy-free agents of 1, as defined in Definition 3. This is because there is a $\frac{1}{n}$ probability of n envy-free agents, and 0 envy-free agents otherwise. By this measure alone, this is equivalent to a mechanism that always has 1 envy-free agent, such as a "random dictator" mechanism. The "random dictator" picks an agent at random and gives that agent its most preferred item along with the maximum share of the divisible resource (i.e. $\max(T, 0)$), with the remaining resources allocated to other agents independently of all agent bids. As the probability of being the dictator does not depend on reported types, no agent can benefit by misreporting its type.

The maximum expected number of envy-free agents is n, and this implies that every outcome is envy-free. However, as shown in the previous subsection, this is not possible for a truthful mechanism.

Theorem 2. *A truthful (in expectation) mechanism for the RA-RD problem with n agents has an expected number of envy-freeness of at most* $(n - 1 + \frac{1}{n})$.

Proof. From Theorem 1, the maximum probability of an envy-free outcome is $\frac{1}{n}$, where there are n envy-free agents. The remaining outcomes, with probability $\frac{n-1}{n}$, can have at most $(n-1)$ envy-free agents. This gives an expected number of envy-free agents of $n\frac{1}{n} + (n-1)\frac{n-1}{n} = n - 1 + \frac{1}{n}$ □

The GPEF was maximised with a fairly simple mechanism, and in the rest of this section we present mechanisms for maximising the expected number of envy-free agents. The first is a mechanism that achieves the bound in Theorem 2 for two agents, followed by a more general mechanism with expected number of envy-free agents of at least $(n-1)$, falling short of the bound by $\frac{1}{n}$.

The 2 Agent Case

For $n = 2$, this bound, $\frac{3}{2}$, can be reached with the following mechanism. Let I_j denote the point of indifference for agent j, which is the division of the divisible resource such that all bundles have equal value. For two agents, this can be represented as a single value, as the divisions must sum to T, and can be calculated as:

$$v_{j,1} + I_j = v_{j,2} + (T - I_j) \Rightarrow \quad I_j = \frac{1}{2}(v_{j,2} - v_{j,1} + T)$$

The mechanism chooses an agent at random, and uses that agent's point of indifference to determine bundles. Agents are then randomly assigned to a bundle. Each agent has a $\frac{1}{2}$ probability of being assigned each indivisible resource, and has an constant expected share of the divisible resource:

$$\bar{r}_1 = \bar{r}_2 = \frac{1}{2}\left(\frac{I_1 + (T - I_1)}{2} + \frac{I_2 + (T - I_2)}{2}\right) = \frac{T}{2}$$

So by Lemma 1, this mechanism is truthful in expectation. The agent chosen to set the bundles will be envy-free with either bundle, while the other agent will prefer one bundle, so there is a probability of $\frac{1}{2}$ this agent will be envious. This gives expected number of envy-free agents of $\frac{3}{2}$ and a GPEF of $\frac{1}{2}$. Thus, based on both measures of envy-freeness, the worst-case behaviour cannot be improved.

The $n > 2$ Agent Case

Our mechanism is a random distribution over deterministic mechanisms that are modifications to a VCG allocation with $(n-1)$ agents, based on the randomised technique proposed by Faltings [2005]. The mechanism proceeds as follows:

1. Find f, the efficient allocation for all agents in N. The value of this efficient allocation is $\bar{C} = \sum_{i \in N} v_i(f)$.
2. Next, randomly select an agent $x \in N$, with equal probability over all agents, as the agent to be "ignored".
3. Find f_{-x} and $f_{-\{i,x\}}$, the efficient allocations for agents $N \setminus \{x\}$ and $N \setminus \{i, x\}$ respectively, for all agents $i \neq x$.
4. Assign non-ignored agents according to f_{-x}, giving agent x the leftover item.
5. Agents make payments according to r_i^x for each agent $i \neq x$, and r_x^x for agent x, as in the following equations.

$$r_i^x = -C^x + v_i(f_{-x}(i)) + C_{-i}^x + \frac{T}{n} - \frac{\bar{C}}{n}, i \neq x \quad (1)$$

$$r_x^x = T - \sum_{i \neq x} r_i^x$$
$$= (n-2)C^x - \sum_{i \neq x} C_{-i}^x + \frac{T}{n} + \frac{(n-1)}{n}\bar{C} \quad (2)$$

Where $C^x = \sum_{j \neq x} v_j(f_{-x}(j))$ is the value of the efficient allocation excluding x, and $C_{-i}^x = \sum_{j \neq \{i,x\}} v_j(f_{-\{i,x\}}(j))$ is the value of the efficient allocation excluding $\{x, i\}$.

The payment for agent x is calculated based on the other agents' payments to ensure strong budget balance, i.e. the sum of all payments is equal to T. The payment r_i^x is made up of three parts. The first three terms in Equation 1 are the VCG payments with Clarke pivot payments in an allocation setting with agent x ignored. For this part of the payment function, along with the allocation function f_{-x}, the agents will have no incentive to misreport. Additionally, VCG mechanisms with Clake pivot payments are known to be envy-free when agents only receive one item [Leonard, 1983; Cohen *et al.*, 2010], so there will be no envy between non-ignored agents. The term $\frac{T}{n}$ is added equally to all agents, so will not affect envy or truthfulness. It is added to ensure payments sum to T. The final term, $\frac{-\bar{C}}{n}$, is added to ensure no agents are envious of the ignored agent. It is added equally to all agents, so will not create envy between non-ignored agents. This breaks the incentive-compatibility of the VCG payments, as it depends on all agents' reported values. When considering *expected* utility, agents have a $\frac{1}{n}$ probability of paying $\frac{(n-1)}{n}\bar{C}$ and an $\frac{(n-1)}{n}$ probability of paying $\frac{-\bar{C}}{n}$, so in expected utility the term cancels out. This means the mechanism remains truthful in expectation. If the value of the efficient allocation is at least T, then all agents will have a non-negative expected utility.

While non-ignored agents are not envious of each other, the pricing must also ensure they are not envious of the ignored agent. Agent i is envious of agent x iff:

$$v_i(f_{-x}(i)) - r_i^x < v_i(f_{-x}(x)) - r_x^x$$
$$\Rightarrow \bar{C} < v_i(f_{-x}(x)) + C_{-i}^x + \sum_{j \neq x} C_{-j}^x - (n-1)C^x \quad (3)$$

Since, assuming non-negative agent values, $C^x \geq C_{-i}^x$, then $\bar{C} \geq \bar{C} + \sum_{i \neq x} C_{-i}^x - (n-1)C^x$. Also, for any agents $\{i, x\}$, we have $\bar{C} \geq C^x \geq C_{-i}^x + v_i(f_{-x}(x))$. Otherwise the efficient allocation used for C^x could have been improved by using allocation $f_{-\{i,x\}}$ and switching agent i to item $f_{-x}(x)$. Thus we have:

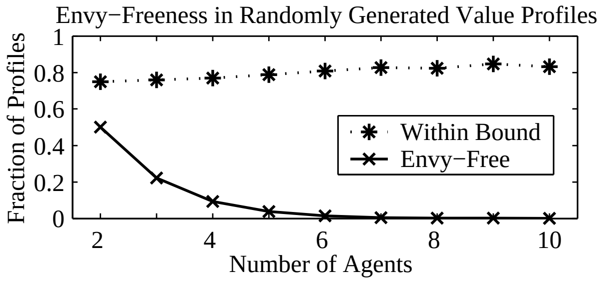

Figure 1: Fraction of value profiles that give outcomes within the worst case bounds, and where all outcomes are envy-free.

$$\bar{C} \geq v_i(f_{-x}(x)) + C_{-i}^x + \sum_{i \neq x} C_{-i}^x - (n-1)C^x$$

As no agent can be envious of the ignored agent, for any choice of x, there will be at least $(n-1)$ envy-free agents. This is the minimum expected number of envy-free agents, but short of the upper bound by $\frac{1}{n}$.

3.3 Empirical tests

The mechanism described in Section 3.2 for $n > 2$ agents does not meet the bound for guaranteed probability of envy-freeness or expected number of envy-free agents. While $(n-1)$ agents are guaranteed to be envy-free, the excluded agent may be envious in all outcomes. We test our mechanism empirically by generating random value profiles, where each agent's value for an item is drawn from a uniform distribution in the range $[0, 1]$. Including negative values did not noticeably affect our results. We then calculated the expected number of envy-free agents and the probability of envy-freeness for each value profile. At least 2500 random value profiles were generated for each n.

The plot in Figure 1 summarises the fraction of value profiles that give at least $\frac{1}{n}$ probability of envy-freeness, and the profiles that always give envy-freeness. For this mechanism, outcomes either have 0 or 1 envious agents, so all outcomes that give a probability of envy-freeness of at least $\frac{1}{n}$ also have an expected number of envy-free agents of at least $(n-1+\frac{1}{n})$. The dotted line in the plot shows that the majority of profiles fall within the optimal bound for these two measures, and this fraction increases with additional agents. However, there is still a significant fraction of profiles for which this mechanism falls short of this bound. So these worst-case profiles are not rare, special cases. The solid line shows that the fraction of ideal cases, where all outcomes are envy-free, for this mechanism rapidly approaches zero. So with this mechanism, an input that will always give an envy-free outcome becomes extremely rare as n increases.

4 Relation to Heterogeneous Item Allocation

This randomised approach to the room assignment-rent division problem, along with the measures used to assess randomised mechanisms can be used in related problems. The problem of budget balanced, efficient allocation involves distributing a set of heterogeneous items to a set of agents such that the items are allocated efficiently, the sum of all agents' payments is zero (strong budget balance), and no agent benefits from misreporting preferences. In variations of this problem, there can be a different number of agents and items, and agents may not necessarily have unit demand. However, due to the Green-Laffont impossibility theorem, there is no efficient mechanism that is DSIC and strong budget balanced.

The RA-RD mechanism for $n > 2$ agents, described in Section 3.2, with the $\bar{C}$ and T terms removed from payment functions is strong budget balanced and DSIC. This is because the VCG mechanism used after an agent is ignored is DSIC and the ignored agent is payed so as to achieve strong budget balance. While not efficient deterministically, the Pareto efficiency of randomised mechanisms can be assessed by measures similar to those used for envy-freeness in RA-RD. For each choice of ignored agent, the remaining $(n-1)$ agents are assigned to an efficient allocation for those agents. Thus, in every outcome, the expected number of agents over which the allocation is efficient is at least $(n-1)$. This is similar to the property of expected number of envy-free agents. Note that this will hold for general allocation settings, not just those where each agent receives at most one item.

In the restricted case where each agent receives at most one item, and where $m \leq n$, for at least one chosen ignored agent the overall allocation for all n agents is efficient. For $n = m$, there is at least one agent who, when ignored, does not change the efficient allocation of the remaining agents. Furthermore, if $m < n$, then ignoring any of the agents that were left unallocated in the efficient allocation will also leave the allocation unchanged. In cases where the allocation is unchanged, then the final outcome will be efficient over all agents. As there are n different outcomes, and $n - m$ agents who receive no item in the efficient allocation, this gives a worst-case probability of an efficient allocation of $\frac{1}{n}$ for cases where $m = n$, or $\frac{n-m}{n}$ for cases where $m < n$. This measurement is analogous to the guaranteed probability of envy-freeness, and from Theorem 1 it is also the best achievable for $n = m$.

5 Conclusions and Future Work

In this work we presented randomised mechanisms for achieving envy-freeness in the room assignment-rent division problem. A deterministic mechanism is unable to provide an envy-free outcome while ensuring agents have no incentive to misreport their preferences. For a randomised mechanism, there are several possible outcomes, so evaluating and comparing these mechanisms by purely deterministic measures is not always suitable. We presented measures of envy-freeness appropriate for comparing randomised mechanisms.

Calculating envy between agents' lotteries of outcomes is not an effective measure in the RA-RD problem, as we show it is trivial to achieve this in mechanisms, and it does not consider the quality of final outcomes. Instead we focused on measuring the GPEF, which shows, in the worst case, what probability the mechanism will achieve the ideal outcome of envy-freeness in all agents. We also propose assessing mechanisms based on the expected number of envy-free agents, which can give an expected level of quality where the ideal outcome is unlikely. For these measures on the RA-RD problem, we provided upper bounds for strategy-proof randomised mechanisms.

These measures can be applied to mechanisms in other problems where truthful, deterministic, envy-free mechanisms are impossible. Similar measures can also be used on other qualities, such as Pareto efficiency. Efficiency cannot be achieved with strong budget balance and incentive compatibility, but a randomised mechanism can guarantee a minimum probability of efficiency in the worst-case.

References

[Abdulkadiroğlu *et al.*, 2004] A. Abdulkadiroğlu, T. Sönmez, and M. U. Ünver. Room assignment-rent division: A market approach. *Social Choice and Welfare*, 22(3):515–538, 2004.

[Alkan *et al.*, 1991] A. Alkan, G. Demange, and D. Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4):1023–39, 1991.

[Alon *et al.*, 2010] N. Alon, F. Fischer, A. D. Procaccia, and M. Tennenholtz. Sum of us: Strategyproof selection from the selectors. Working paper, March 2010.

[Andersson and Svensson, 2008] T. Andersson and L.-G. Svensson. Non-manipulable assignment of individuals to positions revisited. *Mathematical Social Sciences*, 56(3):350–354, 2008.

[Andersson, 2009] T. Andersson. A general strategy-proof fair allocation mechanism revisited. *Economics Bulletin*, 29(3):1717–1722, 2009.

[Brams and Taylor, 1996] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

[Cohen *et al.*, 2010] E. Cohen, M. Feldman, A. Fiat, H. Kaplan, and S. Olonetsky. Truth and envy in capacitated allocation games. Discussion Paper Series, Hebrew University, Jerusalem, 2010.

[Faltings, 2005] B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-Mediated Electronic Commerce VI*, 30–43., 2005.

[Green and Laffont, 1979] J. R. Green and J.-J. Laffont. *Incentives in public decision-making*. North-Holland Publishing Company, Amsterdam, 1979.

[Haake *et al.*, 2002] C.-J. Haake, M. G. Raith, and F. E. Su. Bidding for envy-freeness: A procedural approach to n-player fair-division problems. *Social Choice and Welfare*, 19(4):723–749, 2002.

[Kojima, 2009] F. Kojima. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1):134–142, 2009.

[Leonard, 1983] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–79, 1983.

[Moulin and Bogomolnaia, 2001] H. Moulin and A. Bogomolnaia. A simple random assignment problem with a unique solution. *Economic Theory*, 19(3):623–636, 2001.

[Procaccia, 2010] A. D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proc. of AAAI'10*, Atlanta, GA, 2010.

[Su, 1999] F. E. Su. Rental harmony: Sperner's lemma in fair division. *American Mathematical Monthly*, 106(10):930–942, 1999.

[Sun and Yang, 2003] N. Sun and Z. Yang. A general strategy proof fair allocation mechanism. *Economics Letters*, 81(1):73–79, 2003.