

Online Fair Division: Analysing a Food Bank Problem

Martin Aleksandrov and Haris Aziz and Serge Gaspers and Toby Walsh

NICTA and UNSW, Sydney, Australia

{martin.aleksandrov, haris.aziz, serge.gaspers, toby.walsh}@nicta.com.au

Abstract

We study an online model of fair division designed to capture features of a real world charity problem. We consider two simple mechanisms for this model in which agents simply declare what items they like. We analyse axiomatic properties of these mechanisms such as strategy-proofness and envy-freeness. Finally, we perform a competitive analysis and compute the price of anarchy.

1 Introduction

Resource allocation is a fundamental problem facing society. How do we share scarce and often costly resources between different parties? Due to environmental, economic and technological changes, there is an ever increasing pressure on the allocation of resources. The theoretical foundations of resource allocation have been developed using simple abstract models. For example, one simple model for resource allocation is fair division. Fair division problems are typically categorised along several orthogonal dimensions: divisible or indivisible goods, centralised or decentralised mechanisms, cardinal or ordinal preferences, etc. (e.g. [Aziz *et al.*, 2014; Chevaleyre *et al.*, 2006]). However, such categories do not capture the richness of many real world fair division problems. This has motivated a call to develop more complex and realistic models and mechanisms [Walsh, 2015]. In this paper, we respond to this call by studying mechanisms for an online fair division problem first proposed in [Walsh, 2014].

2 The Food Bank problem

Unfortunately, even in developed countries, poverty remains a serious problem. For example, the 2012 report “*Poverty In Australia*” estimated that over 2 million people (12.5% of the population) are within the official definition of poverty (less than half the median income) [Davidson, 2012]. Amongst the young and old, the statistics are even worse (roughly 1 in 6 children, and 1 in 4 pensioners). These people struggle to feed themselves and increasingly call upon food banks to help. Food Bank Australia sees the demand on their services increase by over 10% per annum. For this reason, they are keen on improving the efficiency of their operations.

In cooperation with a social startup, FoodBank Local, we have been helping Food Bank Australia develop technologies to operate more effectively. So far, this has involved building an app to help collect and deliver donated food. This app uses our vehicle routing solver to route their trucks. We are now turning our attention to how the donated food is allocated to different charities. This is an interesting fair division problem. It has many traditional features. We want to allocate food *fairly* between the different charities that feed different sectors of the community. Goods are mostly *indivisible*. The allocation does not use money as these are all charities. However, the problem also has other features not traditionally found in the academic literature on fair division. One of the main novelties is that it is *online*. Food is donated throughout the day and we must start allocating and distributing it almost immediately, before we know what else will be donated. We have therefore formulated an online model of their fair division problem, and studied mechanisms that can fairly and efficiently allocate the donated food.

3 Online fair division

We have k agents and m items. Each agent has some (private) utility for each item. One of the items appears at each time step, each agent bids (i.e. reports a value) for it and the allocation mechanism must assign it to one of the agents. The next item is then revealed. This continues for m steps. To allocate items in this online model, we consider a simple class of bidding mechanisms in which agents merely declare if they like items or not. For instance, suppose a public kitchen has a rotating schedule of fixed meals. A product that is not in an upcoming meal would not be liked by the corresponding agent. In addition, she would equally like all components of the meal. The LIKE mechanism allocates the next item uniformly at random between agents that bid for the item. An item is possibly allocated to an agent by the LIKE mechanism if an agent bids for it and necessarily allocated if no other agent bids for it.

One drawback of the LIKE mechanism is that agents can get unlucky. It is possible for them to bid for every item but have every coin toss go against them and not be allocated anything at all. This is highly undesirable in our Food Bank setting. A whole sector of the population will then not be fed that night. We therefore consider a slightly more sophisticated mechanism that helps tackle this problem. The

BALANCED LIKE mechanism tries to balance the number of items allocated to agents compared to the LIKE mechanism. It allocates the next item uniformly at random between those agents that value it that have so far received the fewest items. The BALANCED LIKE mechanism is less likely to leave agents empty handed than the LIKE mechanism. In particular, an agent is *guaranteed* to be allocated at least one item for every k items that they bid for. However, there is no guarantee that it *necessarily* returns balanced allocations.

Given the order of items, we can compute the actual outcome of both the LIKE and BALANCED LIKE mechanisms efficiently. Each of the m steps takes $O(k)$ time. Supposing agents bid sincerely, computing the probability an agent gets a particular item, as well as their expected utility is more challenging as there are $O(k^m)$ possible outcomes. With the LIKE mechanism, the probability that agent i gets item j is simply $\frac{1}{q_j}$ where q_j is the number of agents who bid for j . The expected utility is then $\sum_{j=1}^m \frac{u_i(j)}{q_j}$ where $u_i(j)$ is the private utility of agent i for item j . With the BALANCED LIKE mechanism, we can compute the probability that an agent gets a particular item using dynamic programming. This exploits the fact that the mechanism is Markovian. It doesn't care how we get to a particular state, just how many items each agent has at this point. The states represent the number of items allocated to each agent. We can compute the probability that an agent gets a particular item, as well as the expected utility of an agent in $O(m^k)$ space and time. Note that k is typically smaller than m so $O(m^k)$ is likely to be better than $O(k^m)$.

4 Strategy-proofness

As is common in the literature, we will consider the axiomatic properties of these mechanisms. For example, we say that a mechanism for online fair division is *strategy-proof* (weak dominant strategy incentive compatible) if, with knowledge of the items still to be revealed, the order in which they will be revealed, and the private utilities of the other agents, an agent cannot increase her expected utility by bidding differently to their true preferences. We might prefer strategy-proof mechanisms as agents cannot manipulate the outcome to improve their utility at the expense of agents who are less sophisticated or knowledgeable.

Theorem 1 *The LIKE mechanism is strategy-proof.*

With the BALANCED LIKE mechanism, balancing the size of the allocations has an unfortunate side effect: an agent can now manipulate the outcome to increase their expected utility by bidding strategically. In particular, an agent may choose not to bid for an item now in the knowledge that this will bias future allocation rounds in their favour. Such manipulations may decrease the balance of the final allocation.

Theorem 2 *The BALANCED LIKE mechanism is not strategy-proof even when restricted to 0/1 utilities.*

Proof. Suppose we are allocating items a , b and c in this order between agents 1, 2 and 3 with agent 1 having utility 1 for all items, agent 2 for a and c , and agent 3 for b alone. Then bidding sincerely gives agent 1 an expected utility of $\frac{9}{8}$

but this increases to $\frac{5}{4}$ if agent 1 strategically bids only for items b and c and supposing the other agents bid sincerely. \square

It is a strong assumption to suppose that a strategic agent has full knowledge of the items still to be revealed, the order in which they will be revealed, and the private utilities of the agents for these items. In practice, agents may only have partial knowledge. This will greatly limit the willingness of, say, a risk averse agent to be strategic. For instance, if there is a chance that only items that you do not value will arrive in the future, a risk averse agent will always sincerely bid for an item that arrives now which they value. Interestingly, when limited to just two agents and 0/1 utilities, the BALANCED LIKE mechanism becomes strategy-proof even under our strong assumption of complete knowledge.

Theorem 3 *The BALANCED LIKE mechanism is strategy-proof with 2 agents and 0/1 utilities.*

Proof sketch. It is sufficient to prove that sincere play is dominant strategy for one of the agents, say 1. Thus, we show by induction that she does not have an incentive to misreport her sincere valuation of any item.

Given a revealing order of the items, item i is allocated at to either agent 1 or 2, or neither of them. We view the allocation process as an allocation tree as follows. A node labelled by $(i, (x, y))$ encodes a decision point in the allocation process where (1) item i is allocated and (2) x and y denote the numbers of items already allocated to agents 1 and 2, respectively. Depending on the allocation decision that is taken at node $(i, (x, y))$, we arrive at child node of $(i, (x, y))$ which is labelled either $(i + 1, (x + 1, y))$, $(i + 1, (x, y + 1))$ or $(i + 1, (x, y))$ supposing item i is allocated to agent 1, agent 2 or neither of them. Let $T(i, (x, y))$ be the allocation sub-tree starting from a node $(i, (x, y))$ at round i and $U_1(T(i, (x, y)))$ be the total expected utility of agent 1 over that tree supposing agents act sincerely.

Observation 1 *The allocation tree has the following memory-less property: if there is a node labelled by $(i, (x, y))$ and a node labelled by $(i, (x', y'))$ such that $x - y = x' - y'$, then the sub-trees $T(i, (x, y))$ and $T(i, (x', y'))$ are identical.*

The following lemmas can be proved by analysing the allocation tree. The base cases are trivial, with $i = m$. In the step cases, we suppose the lemma statements hold for items $i + 1$ to m and thus we show they hold for items i to m .

Lemma 1 *For any integers x and y , and for all $i = 1, \dots, m$, $U_1(T(i, (x, y))) \geq U_1(T(i, (x - 1, y + 1)))$*

Lemma 2 *For any integers x and y , and for all $i = 1, \dots, m$, $U_1(T(i, (x, y))) \geq U_1(T(i, (x - 1, y)))$*

Lemma 3 *For any integers x and y , and for all $i = 1, \dots, m$, $U_1(T(i, (x, y))) \geq U_1(T(i, (x, y + 1)))$*

Using these lemmas, we can prove that agent 1 has no incentive to misreport her sincere valuation of any item. Let u'_1 denotes agent 1's insincere bid, i the last item for which agent 1 does not bid sincerely, and u_1 her sincere valuation of item i . Suppose that we are at node $(i, (x, y))$ and let us consider 2 cases. First, let $u'_1(i) = 1$ whilst $u_1(i) = 0$. By Observation 1 and Lemma 3, agent 1 has no incentive to bid $u'_1(i)$ for item i . Second, let $u'_1(i) = 0$ whilst $u_1(i) = 1$. We focus on the

node $(i, (x, y))$ which leads to different sub-trees depending on whether agent 1 reports u_1 or u'_1 .

- If agent 2 does not bid for i , then agent 1 certainly gets i under u_1 but no one gets i if agent 1 reports u'_1 . Under u_1 we arrive at node $(i+1, (x+1, y))$, whereas under u'_1 we arrive at node $(i+1, (x, y))$. By Lemma 2, u_1 yields at least as much utility as u'_1 since $U_1(T(i+1, (x+1, y))) \geq U_1(T(i+1, (x, y)))$.
- If agent 2 bids for i and $x < y$, then agent 1 receives item i under u_1 and agent 2 receives the item under u'_1 . Since, by Lemma 1, $U_1(T(i+1, (x+1, y))) \geq U_1(T(i, (x, y+1)))$, reporting u_1 yields at least as much utility to agent 1 as u'_1 .
- If agent 2 bids for i and $x = y$, then under u_1 there are two children of $(i, (x, y))$ labelled $(i+1, (x+1, y))$ and $(i+1, (x, y+1))$, respectively. But under u'_1 , there is only one child of $(i, (x, y))$, i.e. $(i+1, (x, y+1))$. By Lemma 1, $U_1(T(i+1, (x+1, y))) \geq U_1(T(i+1, (x, y+1)))$ and thus u_1 yields at least as much utility to agent 1 as u'_1 .
- If agent 2 bids for i and $x > y$, then agent 2 receives item i , no matter how agent 1 bids. Therefore, agent 1 has no incentive to report u'_1 rather than u_1 .

This completes the proof of the theorem as agent 1 does not have incentive to play insincerely in all cases. \square

The latter proof critically requires 0/1 utilities. It is easy to give examples with more general utilities where the BALANCED LIKE mechanism is not strategy-proof even with 2 agents.

Example 1 Consider 2 agents and 2 items, a and b . Agent 1 has utility $\frac{1}{2}$ for both items, and agent 2 has utility $\frac{1}{4}$ for item a and $\frac{3}{4}$ for item b , normalized to sum up to 1. If agents bid sincerely for both items, then agent 2 has an expected utility of $\frac{1}{2}$. However, by bidding strategically only for item b , agent 2 can increase their expected utility to $\frac{3}{4}$.

5 Impact on welfare

Strategic play can have both positive or negative effect on the welfare of the community. We measure this from utilitarian and egalitarian perspective. The former is the sum of the expected utilities of the agents and thus measures the collective welfare. The latter is the expected utility of the worst-off agent and thus measures individual welfare. As these two objectives may conflict, we consider both measures. We consider pure Nash equilibria in which no agent can get strictly greater expected utility by changing their strategy. However, there are pure Nash equilibria that have much smaller egalitarian and utilitarian welfare than sincere play for both mechanisms.

Theorem 4 *There are instances with 0/1 utilities and k agents, where the egalitarian and utilitarian welfare of sincere play in the LIKE and BALANCED LIKE is k times the corresponding welfare of at least one pure Nash equilibrium.*

Proof. Consider an instance with k agents and k items. For each $i \in \{1, \dots, k\}$, agent i values item i and no other item.

For sincere play, item i is assigned to agent i in both the LIKE and BALANCED LIKE mechanisms, giving an egalitarian utility of 1 and a utilitarian utility of k . Let us now consider the pure Nash equilibrium where each agent bids for all items. In the LIKE mechanism, with these bids, each agent is allocated each item with probability $\frac{1}{k}$. Since each agent values exactly one item, this gives egalitarian welfare of $\frac{1}{k}$ and utilitarian welfare of 1. In the BALANCED LIKE mechanism, each agent is allocated exactly one item. The probability that this item is the one she likes is $\frac{1}{k}$, giving again egalitarian welfare of $\frac{1}{k}$ and utilitarian welfare of 1. \square

For the LIKE mechanism, a pure Nash equilibrium cannot lead to greater egalitarian or utilitarian welfare than sincere play as no player has an incentive not to bid for an item she likes. Also, the example in the last proof involves many agents that bid for items for which they have no value. As a result, they end up again with lower expected utility. Therefore, we further consider a subset of pure Nash equilibria by supposing a small utility cost to liking (or taking delivery of) an item. We call these *simple* pure Nash equilibria. Note that sincere play is the only simple pure Nash equilibrium for the LIKE mechanism, and therefore, there is no difference in welfare between sincere play and simple pure Nash equilibria.

For the BALANCED LIKE mechanism, simple pure Nash equilibria have the same utilitarian welfare as sincere play, as each item is assigned to an agent who likes it. However, we next show that a simple pure Nash equilibrium may have smaller or greater egalitarian welfare than sincere play.

Theorem 5 *There are instances with 0/1 utilities where the egalitarian welfare of sincere play in the BALANCED LIKE mechanism is strictly greater than the egalitarian welfare of each simple pure Nash equilibrium.*

Proof. The proof of Theorem 2 gives an instance where the unique simple pure Nash equilibrium has less expected egalitarian utility than sincere play. \square

Theorem 6 *There are instances with 0/1 utilities where the egalitarian welfare of sincere play in the BALANCED LIKE mechanism is strictly smaller than the egalitarian welfare of each simple pure Nash equilibrium.*

Proof. Consider the fair division of items a to f in alphabetical order between agents 1, 2 and 3 with the following preferences: agent 1 has utility 1 for items a to c and 0 for d to f , agent 2 has utility 1 for a, c, e and f and 0 for b and d and agent 3 has utility 1 for a, b, d and f and 0 for c and e . By running the BALANCED LIKE mechanism, one always obtains an allocation with egalitarian welfare of 1, except when the items are allocated to the agents according to the sequence of agents $(2, 1, 1, 3, 2, 3)$, in which case the egalitarian welfare is 2. By analysing the allocation tree of the BALANCED LIKE mechanism, one can see that this instance has a unique simple pure Nash equilibrium which favours this allocation and in which agent 1 does not bid for item a and all other bids are the same. We obtain an egalitarian welfare of $\frac{13}{12}$ for sincere play and $\frac{9}{8}$ for the simple pure Nash equilibrium. \square

6 Envy-freeness

How fair are these mechanisms? Is the BALANCED LIKE mechanism more fair in some sense than the LIKE mechanism. Since the outcomes of our mechanisms are random, we consider fairness notions both *ex post* (with respect to the actual allocation achieved in a particular world) and *ex ante* (with respect to the expected utility over all possible worlds). One notion of fairness commonly considered in the fair division literature is envy freeness [Brams and Taylor, 1996]. An agent *envies ex post/ex ante* another agent if their utility/expected utility of the other agent's allocation is greater than their utility/expected utility of their own allocation. A mechanism is *envy free ex post/ex ante* if no agent envies another *ex post/ex ante*. We also consider a weaker notion. An agent has *bounded envy ex post/ex ante* of another agent if there exists a constant r such that in every case their utility/expected utility of the other agent's allocation is at most r greater than their utility/expected utility of their own allocation. Similarly, we say that a mechanism is *bounded envy free ex post/ex ante* with constant r if each agent has bounded envy *ex post/ex ante* of every other agent with constant r .

If a mechanism is envy free *ex post/ex ante* then it is bounded envy free *ex post/ex ante*, whilst if a mechanism is (bounded) envy free *ex post* then it is (bounded) envy free *ex ante*. It is easy to show that no mechanism for indivisible items that allocates all items can be envy free *ex post*: suppose we have one indivisible item and two or more agents who bid for it. Regarding the other envy free properties, we prove the following results.

Theorem 7 *Supposing agents act sincerely, the LIKE mechanism is envy free ex ante. It is not bounded envy free ex post, even with 0/1 utilities and 2 agents.*

Proof. To prove envy freeness *ex ante*, we perform induction over the number of items. In the base case, we have no items to allocate, each agent receives an expected utility of 0, and no agent envies another *ex ante*. For the induction step, we suppose the allocation of the first $m - 1$ items is envy free *ex ante*, and consider the m th item which is allocated. Suppose j ($\leq k$) agents have non-zero utility for the m th item. Then each agent receives this item in $\frac{1}{j}$ of the possible worlds. This means that the new allocation remains envy free *ex ante*.

To show that the LIKE mechanism is not bounded envy free *ex post* even with 0/1 utilities, suppose 2 agents have utility 1 for all m items. There is one outcome in which the first agent gets lucky and is assigned every item. However, in this case, the other agent assigns utility of m units greater to the first agent's allocation than to their own (empty) allocation. \square

As the LIKE mechanism is strategy-proof, it seems reasonable to suppose agents act sincerely. By comparison, when limited to 0/1 utilities, the BALANCED LIKE mechanism is both envy free *ex ante*, and bounded envy free *ex post*.

Theorem 8 *Supposing agents act sincerely, the BALANCED LIKE mechanism with 0/1 utilities is envy free ex ante and bounded envy free ex post with constant 1.*

Proof sketch. Both proofs use induction on the number of items. For envy freeness *ex ante*, the induction step uses case analysis to show that the expected increase in utility for an

agent is at least as large as their expected increase in utility for the allocation of any other agent. For bounded envy freeness *ex post*, the induction step again uses case analysis to show that the envy is at most 1 unit. \square

It is not hard to show that with general utilities, the BALANCED LIKE mechanism is no longer envy free *ex ante*, or bounded envy free *ex post* (or even, *ex ante*). Balancing the allocation of items may prevent an agent who values an item greatly from being allocated it.

Example 2 *Consider 2 agents and 2 items, a and b . Suppose agent 1 has utility 0 for a and p for b , but agent 2 has utility 1 for item a and $p - 1$ for item b where $p > 2$. Note that both agents have the same sum of utilities for the two items. If agents bid sincerely then agent 2 gets an expected utility of just 1 and envies *ex ante* agent 1's allocation which gives agent 2 an expected utility of $p - 1$. As p is unbounded, agent 2 does not have bounded envy *ex post* or *ex ante* of agent 1.*

To conclude, on the basis of envy freeness, provided utilities are 0/1 (or close to this), we might consider the BALANCED LIKE mechanism to be somewhat more fair than the LIKE mechanism. On the other hand, when utilities are not only 0/1 (or close to this), we might consider the BALANCED LIKE mechanism to be somewhat less fair than the LIKE mechanism.

7 Competitive analysis

A powerful technique to study online mechanisms is competitive analysis [Sleator and Tarjan, 1985]. This identifies the loss in efficiency due to the data arriving in an online fashion. We say that a randomized mechanism M for online fair division is c -competitive from an egalitarian/utilitarian perspective if there exists a constant a such that whatever the input sequence of items π :

$$SW_{OPT} \leq c \cdot SW_M(\pi) + a$$

where $SW_M(\pi)$ is the egalitarian/utilitarian social welfare of the mechanism on π , and SW_{OPT} is the optimal egalitarian/utilitarian social welfare of an (offline) assignment. We suppose agents bid sincerely. The following results hold irrespective of the model of the adversary (oblivious, or adaptive offline).

The LIKE mechanism is competitive when the number of agents is bounded, even with general utilities.

Theorem 9 *With general utilities and k agents, the LIKE mechanism is k -competitive from an egalitarian or utilitarian perspective.*

Proof. With the LIKE mechanism, the worst case for every agent is that every other agent bids against them. Hence, the worst case is that their expected social welfare is $\frac{1}{k}$ the smallest sum of utilities. By comparison, the best case for an agent is that they receive the sum of their utilities. Hence, the competitive ratio from an egalitarian or utilitarian perspective is at worst k .

From an egalitarian perspective, this bound is met even when utilities are just 0 or 1. Consider k^2 items being divided between k agents. The first agent has utility of 1 for the first k items and 0 for all remaining items. The other agents have

utility 1 for all items. The optimal offline allocation achieves egalitarian social welfare of k units, but the egalitarian social welfare of the LIKE mechanism is just 1 unit.

From a utilitarian perspective, this bound is met even with just k items. Suppose the i th agent has an utility of $1 - (k-1)\epsilon$ for the i th item, and ϵ for all other items where ϵ is a small non-zero constant. Note that the sum of the utilities for any agent is normalized to 1 unit. The optimal utilitarian offline allocation has a social welfare of k units as ϵ goes to zero, whilst the utilitarian social welfare of the LIKE mechanism is just 1 unit. \square

For example, with 2 agents and general utilities, the LIKE mechanism is 2-competitive. That is, the egalitarian or utilitarian social welfare is at least 50% of the optimal (offline) allocation. On the other hand, the BALANCED LIKE mechanism is not competitive even with just 2 agents.

Theorem 10 *With general utilities and 2 agents, the BALANCED LIKE mechanism is not c -competitive from an egalitarian or utilitarian perspective for any constant c .*

Proof. Consider the fair division of items a to d in alphabetical order between agents 1 and 2 with the following preferences: agent 1 has utility ϵ for items a and d , $1 - 2\epsilon$ for b and 0 for c and agent 2 has utility ϵ for b and c , $1 - 2\epsilon$ for d and 0 for a , where $\epsilon > 0$ is a small positive constant. The optimal egalitarian (utilitarian) offline allocation gives items a and b to agent 1 and items c and d to agent 2. This has an egalitarian (utilitarian) social welfare of $1 - \epsilon$ unit ($2 - 2\epsilon$ units). On the other hand, the BALANCED LIKE mechanism results in an egalitarian (utilitarian) social welfare of just 2ϵ (4ϵ), allocating items a and d to agent 1 and the other b and c to agent 2. \square

Finally, when restricted to 0/1 utilities, every allocation of the LIKE or BALANCED LIKE mechanism achieves the utilitarian social welfare of the optimal offline allocation. This is because items only go to agents that value them.

8 Price of anarchy

The price of anarchy is closely related to the competitive ratio but also takes into account agents acting strategically. The price of anarchy measures how the efficiency of a decentralized system degrades due to selfish behavior of its agents compared to imposing a centralized solution based on sincere preferences [Koutsoupias and Papadimitriou, 1999]. From an egalitarian (utilitarian) perspective, the price of anarchy of an online fair division mechanism is the ratio between the optimal egalitarian (utilitarian) social welfare, and the smallest egalitarian (utilitarian) social welfare of any equilibrium strategy. We consider simple pure Nash equilibria (defined in Section 5).

Theorem 11 *With general utilities and k agents, the price of anarchy of the LIKE mechanism is k for egalitarian welfare, and for utilitarian welfare is greater than $k - \epsilon$ for any $\epsilon > 0$.*

Proof. Let us consider the equilibrium strategy with least egalitarian (utilitarian) social welfare. Suppose an agent bids for an item with non-zero utility. The worst case is when every other agent bids against them. This gives an expected utility

which is $\frac{1}{k}$ of the sum of their utilities. By comparison, the best case is that they receive the sum of their utilities.

From an egalitarian perspective, this bound is achieved when k^2 items are divided between k agents, the first agent has utility 1 for the first k items, zero for the rest, and every other agent has utility 1 for every item. Then it is a dominant strategy for the first agent to bid for the first k items, and for all other agents to bid for every item. This gives egalitarian social welfare of 1, compared to the optimal egalitarian social welfare of k units.

From an utilitarian perspective, select ϵ' such that $0 < \epsilon' < \frac{\epsilon}{k \cdot (k-1)}$. The bound is achieved when k items are divided between k agents, the i th agent has utility $1 - (k-1)\epsilon'$ for the i th item and ϵ' for the rest. The dominant strategy is for every agent to bid for every item. In this case, the optimal utilitarian social welfare is $k \cdot (1 - (k-1) \cdot \epsilon') > k \cdot (1 - \frac{(k-1) \cdot \epsilon}{k \cdot (k-1)}) = k - \epsilon$ whilst the utilitarian welfare of the LIKE mechanism is 1. \square

For the BALANCED LIKE mechanism, we have the following lower bounds on the price of anarchy.

Theorem 12 *With 0/1 utilities and k agents, the price of anarchy of the BALANCED LIKE mechanism from an egalitarian perspective is at least k .*

Proof. Consider k^2 items being divided between k agents. The first agent has utility 1 for the first k items and 0 for all remaining items. The other agents have utility 1 for all items. The optimal egalitarian offline allocation gives the first k items to the first agent, and k of the other items to each of the other agents. This has an egalitarian social welfare of k units. On the other hand, a dominant strategy with the BALANCED LIKE mechanism is sincerity. This gives egalitarian social welfare of 1. \square

Theorem 13 *With general utilities and k agents, the price of anarchy of the BALANCED LIKE mechanism from a utilitarian perspective is greater than $k - \epsilon$ for any $\epsilon > 0$.*

Proof. Consider an instance with k items. Select ϵ' such that $0 < \epsilon' < \frac{\epsilon}{k \cdot (k-1)}$. For each $i \in \{1, \dots, k\}$, agent i has utility $1 - (k-1) \cdot \epsilon'$ for item i and utility ϵ' for all other items. In the BALANCED LIKE mechanism, sincere play is the dominant strategy, allocating one item to each agent. The probability that agent i receives item i is $\frac{k-1}{k} \cdot \frac{k-2}{k-1} \cdot \dots \cdot \frac{1}{k-i+1} = 1/k$. Thus, the utilitarian welfare is $1 - (k-1) \cdot \epsilon' + (k-1) \cdot \epsilon' = 1$. The optimal assignment achieves utilitarian welfare of $k \cdot (1 - (k-1) \cdot \epsilon') > k \cdot (1 - \frac{(k-1) \cdot \epsilon}{k \cdot (k-1)}) = k - \epsilon$. \square

Finally, with 0/1 utilities and either mechanism, it is a dominant strategy for agents only to bid for (a subset of) the items for which they have utility. Hence, both mechanisms achieve the optimal utilitarian social welfare. Thus, there is no price of anarchy from an utilitarian perspective in these cases.

9 Experiments

To determine the impact on social welfare of these mechanisms we ran a number of experiments. We used a wide range of problem instances: random 0/1 utilities, random Borda utilities, correlated 0/1 and Borda utilities generated with the Pólya-Eggenberger urn model, as well as 0/1 and Borda utilities from PrefLib.org [Mattei and Walsh, 2013]. For reasons

of space, we report here just results with random 0/1 utilities. We observed similar trends with the other classes (see [Aleksandrov *et al.*, 2015]).

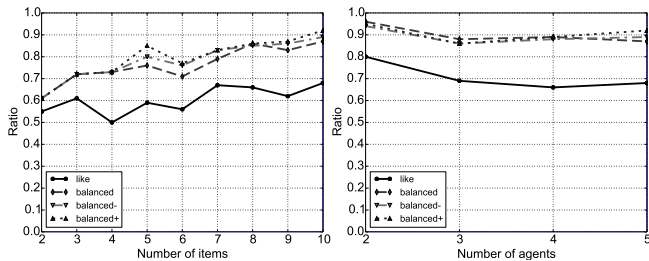


Figure 1: Egalitarian price of anarchy, and competitive ratio of BALANCED LIKE and LIKE mechanisms. (left) varying items for 5 agents, (right) varying agents for 10 items.

We varied the number of agents from 2 to 5, and the number of items from 2 to 10. We sampled 100 instances at each data point, computing the optimal (offline) allocation, and all simple pure Nash equilibria by brute force. In Figure 1, we plot (1) the competitive ratios (“like” and “balanced”), (2) the price of anarchy (“balanced-”) and (3) the ratio between the egalitarian welfare of the best simple pure Nash equilibrium and the optimal allocation (“balanced+”). As these are ratios, we plot geometric means. Arithmetic means are similar. We note that the BALANCED LIKE mechanism (“balanced”) improves the egalitarian welfare compared to the LIKE mechanism (“like”) supposing sincere or strategic play of the agents. Indeed, strategic play of the agents often increases social welfare even in the worst case (“balanced-” compared to “balanced-”), though the effect is small.

In conclusion, BALANCED LIKE performed better than LIKE with 0/1 utilities. Moreover, it remained superior in all our experiments to the LIKE mechanism.

10 Related work

There is a large literature on the fair division of divisible and indivisible goods. Almost all studies assume that all the goods are present initially. There are, however, a few exceptions. Walsh [2011] has proposed an online model of cake cutting. However, in this model the agents arrive over time (not the items), and the goods are divisible (not indivisible). Kash, Procaccia and Shah [2014] have proposed a related model in which agents again arrive over time, but there are now *multiple, homogeneous divisible* goods (and not multiple, heterogeneous indivisible goods as here). Bounded envy freeness is closely related to the “single-unit utility difference” property that Budish, Che, Kojima and Milgrom [2013] prove can be achieved in *offline* fair division with any randomized allocation mechanism that is envy free ex ante.

The LIKE and BALANCED LIKE mechanisms take an item-centric view of allocation. They iterate over the items, allocating them in turn to agents. By comparison, there are agent-centric mechanisms like the sequential allocation procedure which iterate over the agents, allocating items to them in turn [Brams and Taylor, 1999]. These mechanisms have

attracted considerable attention in the AI literature recently (e.g. [Bouveret and Lang, 2011; Kalinowski *et al.*, 2013a; 2013b]). As our matching problem is one-sided (agents have preferences over items, but not vice-versa), we cannot immediately map results from there to here. There are also randomized mechanisms like in [Zhou, 1990] and [Bogomolnaia and Moulin, 2001] which again take an agent-centric view of allocation. It would be interesting future work to consider how such agent-centric mechanisms could be modified to work with online fair division problems.

11 Conclusions

Motivated by our work with a local Food Bank charity, we have studied a simple online model of fair division, as well as two simple mechanisms for this problem. To help decide what mechanism to use in practice, we have studied axiomatic properties of these mechanisms like strategy-proofness and envy-freeness. In addition, we have undertaken a competitive analysis, and computed their price of anarchy. A summary of our results is given in Table 1.

mechanism	LIKE	BALANCED LIKE
strategy-proof	✓	×, ✓ for $k=2$ & 0/1 utilities
envy free (ex ante)	✓	×, ✓ for 0/1 utilities
bound envy free (ex post)	×	×, ✓ for 0/1 utilities
competitive	✓	×
price of anarchy (e)	k	$\geq k$
price of anarchy (u)	$k, 1$ for 0/1 utilities	$\geq k, 1$ for 0/1 utilities

Table 1: Overview of results for k agents. (e) = egalitarian, (u) = utilitarian.

One possible take home message from this table is that we might consider the BALANCED LIKE mechanism if the items can be packaged together so that agents have similar (i.e. 0/1) utility for all packages, and that we should otherwise prefer the LIKE mechanism when this is not possible. In future work, we plan to take into account other important features of this real world allocation problem. For example, as the charities have different abilities to feed their clients, we need a model of online fair division in which the agents have different entitlements. We will need to consider the impact this has on axiomatic properties like strategy-proofness and fairness. We will then be in a position to implement and field a mechanism for online fair division in practice.

Acknowledgments

NICTA is supported by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

Serge Gaspers is the recipient of an Australian Research Council Discovery Early Career Researcher Award (project number DE120101761) and a Future Fellowship (project number FT140100048).

Toby Walsh is supported by the Federal Ministry for Education and Research through the Alexander von Humboldt Foundation, as well as by AOARD Grant FA2386-12-1-4056.

References

- [Aleksandrov *et al.*, 2015] M. Aleksandrov, S. Gaspers, and T. Walsh. Empirical analysis of a food bank problem. At 2nd Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE), 14th AAMAS Conference, 2015.
- [Aziz *et al.*, 2014] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. In *Proceedings of the 13th AAMAS Conference*, pages 1305–1312, 2014.
- [Bogomolnaia and Moulin, 2001] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, October 2001.
- [Bouveret and Lang, 2011] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In Toby Walsh, editor, *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pages 73–78. IJCAI/AAAI, 2011.
- [Brams and Taylor, 1996] S. J. Brams and A. D. Taylor. *Fair division: from cake-cutting to dispute resolution*. Cambridge University Press, Cambridge, 1996.
- [Brams and Taylor, 1999] S. J. Brams and A. D. Taylor. *The win-win solution: guaranteeing fair shares to everybody*. Norton, New York and London, 1999.
- [Budish *et al.*, 2013] E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.
- [Chevaleyre *et al.*, 2006] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J.A. Padget, S. Phelps, J.A. Rodriguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica (Slovenia)*, 30(1):3–31, 2006.
- [Davidson, 2012] P. Davidson. *Poverty in Australia*. 2012. ISBN 978085871082.
- [Kalinowski *et al.*, 2013a] T. Kalinowski, N. Narodytska and T. Walsh. A social welfare optimal sequential allocation procedure. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-2013)*, pages 227–233. AAAI Press, 2013.
- [Kalinowski *et al.*, 2013b] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*, pages 452–458. AAAI Press, 2013.
- [Kash *et al.*, 2014] I. Kash, A. D. Procaccia, and N. Shah. No agent left behind: dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51:579–603, 2014.
- [Koutsoupias and Papadimitriou, 1999] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In C. Meinel and S. Tison, editors, *Proceedings of 16th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413. Springer, 1999.
- [Mattei and Walsh, 2013] N. Mattei and T. Walsh. PrefLib: A library of preference data. In *Algorithmic Decision Theory, Third International Conference, (ADT 2013)*, Lecture Notes in Artificial Intelligence, pages 259–270. Springer, 2013.
- [Sleator and Tarjan, 1985] D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, February 1985.
- [Walsh, 2011] T. Walsh. Online cake cutting. In *Algorithmic Decision Theory, Second International Conference, ADT 2011. Proceedings*, Lecture Notes in Computer Science, pages 292–305. Springer, 2011.
- [Walsh, 2014] T. Walsh. Allocation in practice. In *Proceedings of 37th German Conference on Artificial Intelligence (KI-2014)*, Lecture Notes in Artificial Intelligence, pages 13–24. Springer, 2014.
- [Walsh, 2015] T. Walsh. Challenges in resource and cost allocation. In *Proceedings of the 29th AAAI Conference on AI*, pages 4073–4077. Association for Advancement of AI, 2015.
- [Zhou, 1990] L. Zhou. On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123 – 135, 1990.