Empirical analysis of a Food Bank problem

Martin Aleksandrov
UNSW and NICTA
Kensington 2033
Sydney, Australia

Serge Gaspers
UNSW and NICTA
Kensington 2033
Sydney, Australia

Toby Walsh
NICTA and UNSW
Kensington 2033
Sydney, Australia

ABSTRACT

Hunger is a major problem worldwide. To help alleviate hunger, we are studying mechanisms to distribute donated food more effectively. This is a significant challenge faced by food banks. We study a simple setting capturing features of a real-world charity problem. A number of charities (or agents) bid for a number of food items in an online manner and a mechanism allocates each item to a charity. We focus here on an empirical analysis of the properties of two simple mechanisms for this problem. The main conclusion is that the more sophisticated of the two mechanisms achieves fairer online allocations from utilitarian/egalitarian perspective.

Keywords

Online Fair Division, Optimisation, Multiagent Systems

1. INTRODUCTION

The underpinnings of resource allocation in theory are usually studied through the use of simple models. One such long-standing abstraction is fair division (e.g. [13]). Its instances are characterized based on several orthogonal assumptions: (1) divisible or indivisible goods, (2) centralized or decentralized control, (3) cardinal or ordinal preferences, etc. (e.g. [6]). But, how do we allocate scarce and, often, limited resources in practice? In the real world, we are uncertain about every bit of information that can be used when making such important decisions. For example, the goods may not be all available a priori or the preferences might be incomplete. As a result, any *online* allocation of our resource is expected to be less fair than if we were allocating it in an *offline* manner. This motivates the development of more complex and sophisticated mechanisms for online fair division (e.g [16]). Section 2 introduces the problem setting and two such mechanisms that allocate items to agents, first mentioned in [15]. [1] studies strategy-proofness and envy-freeness of these mechanisms, summarized in Sections 3 and 4. We also add on their proportionality in Section 5. Further, Section 6 discusses their competitiveness and Section 7 reports on our experimental results. We conclude that one of the mechanisms competes with the optimal (offline) mechanism that assumes all data is available a priori.

Appears at: 2nd Workshop on Exploring Beyond the Worst Case in Computational Social Choice. Held as part of the 14th International Conference on Autonomous Agents and Multiagent Systems. May 4th, 2015. Istanbul, Turkey.

2. THE SETTING

It is a misfortune that nowadays many people around the world still live in poverty. This is a significant problem even in developed countries such as Australia. According to the 2012 report "*Poverty In Australia*" (i.e. [7]), approximately 10% of the population cannot afford to put food on their table. As a response to that, they urgently call on the food banks for help. The bank itself caters nearly 90000 meals every day and still cannot meet the increasing demand. For this purpose, the charities are keen on improving their operations and thus distributing the food more efficiently.

We have been working with a social startup, FoodBank Local, towards developing technologies to support Australian charities and thus improving the efficiency of their operations. Thus far, this has involved building an app that helps collecting and delivering donated food. This app uses our vehicle routing solver to route their fleet. However, all this food must be allocated to different charities before being catered to local restaurants and public kitchens. We are therefore turning our attention to how to allocate the donations. This is an interesting and non-trivial fair division problem that combines many traditional and novel features. On the one hand, we want to allocate food *fairly* between the different charities as they feed different sectors of the community. The goods are assumed to be packed and hence *indivisible*. On the other hand, the problem is *online*. All donations arrive throughout the day and we must start allocating and distributing them almost immediately, possibly expecting more donations later on. We have therefore formulated an online model of this fair division problem, and studied mechanisms that can fairly and efficiently allocate the donated food.

In our setting, we suppose there are k agents and m items. Each agent has some (private) utility for each item. One item appears at each time step, 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 how much they like each item. We say that an agent *bids* for an item if they declare that they like the item. For instance, the LIKE mechanism allocates the next item uniformly at random between agents that bid for the item. An allocation is a possible outcome of the LIKE mechanism if each item is given to an agent that bids for it, whilst an allocation is a necessary outcome if no two agents bid for the same item, and each item is given to the agent that bids for it, or to no one if no agent bids for it.

One problem with 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 bid for the item and 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 (i.e. equitable) allocations.

3. STRATEGY-PROOFNESS

We say that an agent bids *sincerely* for an item if they report their private utility for this item. Otherwise, the agent bids *insincerely* for the item. Thus, a mechanism for online fair division is *strategy-proof* if, ($\star$) 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 their expected utility by bidding insincerely. Hence, given a mechanism that is strategy proof, no agent can manipulate the outcome and improve their expected utility at the expense of agents who are sincere in their play.

When using the LIKE mechanism, each agent bidding for a given item gets an equal chance of receiving it. As a result, no agent has an incentive to bid for items that they do not like. This will not increase their expected utility. Furthermore, each agent will always bid for items that they like, as otherwise their expected utility will decrease. Consequently, an agent's best play is to bid sincerely and therefore the LIKE mechanism is strategy-proof. In other words, no agent can manipulate the outcome.

This is however not the case when using BALANCED LIKE even when restricted to 0/1 utilities. Now, an agent may choose not to bid sincerely for a particular item that they like and thus expect to be allocated more items in the following rounds. Such manipulations may decrease the equitability of the final allocation. On the contrary, no agent has incentive to bid for an item that they do not like. Such a strategy will only bias the allocation of the following items in favour of the other agents.

EXAMPLE 1. *Suppose we are allocating items a, b and c in alphabetical order between agents* 1, 2 *and* 3, *with the following utilities. Figure 1 depicts all possible allocations of items a, b and c to agents 1, 2 and 3, that are based on these utilities and can be obtained using the* BALANCED LIKE *mechanism. We consider 2 cases.*

	a	b	c
1	1	1	1
2	0	1	0
3	1	0	1

First, by bidding sincerely, agent 1 *receives an expected utility of* $\frac{9}{8}$. *Second, this value increases to* $\frac{5}{4}$ *supposing she bids strategically only for items b and c, and the other agents bid sincerely. The latter is a strict improvement for agent 1. In fact, no other agent can improve their outcome in this setting.*

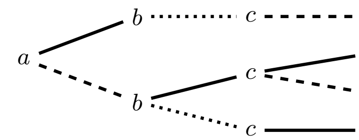

Figure 1: 1-solid, 2-dotted, 3-dashed

The latter result adapts easily to the case with more general utilities. To see that, suppose that agent 1 in Example 1 assigns some strictly greater utility to item c. In this case, her best play is again not to bid for item a. However, note that in this more general setting, the BALANCED LIKE mechanism does not take into account the absolute value of the reported bids when allocating the items. Instead, it only considers whether these bids are equal to zero or not. Hence, no agent has incentive to bid insincerely some non-zero value for an item that they like. Also, no agent has incentive to bid for items that they do not like. As a conclusion, each agent may only have incentive not to bid for items that they like.

Inspired by these arguments, we will consider a subset of pure Nash equilibria by assigning a small utility cost to taking delivery of (liking) an item. Each profile in this subset is called a *simple* pure Nash equilibrium as each agent now may not bid for items that they like. We are interested in how the strategic play of the agents affects the happiness of the community. Two widely accepted aggregate indicators measuring this happiness are utilitarian and egalitarian welfares. The former one is defined as the sum of the expected utilities of the agents whereas the latter one is equal to the minimum expected utility an agent receives. Now, the 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 the simple pure Nash equilibrium. By comparison, the sincerity may not be the best play for an agent when using the BALANCED LIKE mechanism. As a positive consequence, each simple pure Nash equilibrium maximizes the utilitarian welfare. The reason for that is simple - each agent gets items that they like in each possible allocation. Whilst, the strategic play in this case may have either positive or negative effect on the egalitarian welfare. Example 1 indicates that this welfare may decrease. On the other hand, there are instances in which the strategic play of the agents improves it.

EXAMPLE 2. *Suppose the fair division of 6 items in alphabetical order between agents 1, 2 and 3, with the following utilities. Running the* BALANCED LIKE *mechanism, one always obtains an allocation with egalitarian welfare* 1, *except when the items are allocated to the agents according to the sequence of agents*

	a	b	c	d	e	f
1	1	1	1	0	0	0
2	1	1	0	0	1	1
3	1	0	1	1	0	1

$(3, 1, 1, 3, 2, 2)$, *in which case the egalitarian welfare is* 2. *By analysing the allocation tree that is based on the utilities of the agents, one can see that the instance has a unique simple pure Nash equilibrium, which favours this allocation and in which agent* 1 *does not bid for item a. With this strategic move, her expected utility increases to* $\frac{9}{8}$ *from* $\frac{13}{12}$. *The latter is also the value of the egalitarian welfare in the corresponding cases.*

From now onwards, we consider only simple pure Nash equilibria. To conclude, note that ($\star$) is a strong assumption, supposing 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 they do not like will arrive in the future, a risk averse agent will always bid sincerely for an item that arrives now which they like.

4. ENVY-FREENESS

Fairness is important feature that describes the equity of the allocation and the prosperity of the society. But, how fair are our 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 common notion of fairness is envy-freeness (e.g. [3]). 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 allocation. A mechanism is thus *envy free ex post/ex ante* if no agent envies another ex post/ex ante. We also consider a weaker notion of envy-freeness supposing the agents can envy each other, but in a bounded sense. An agent has *bounded envy ex post/ex ante* with constant a of another agent if in every case their utility/expected utility of the other agent's allocation is at most a greater than their utility/expected utility of their own allocation. We say that a mechanism is *bounded envy free ex post/ex ante* with constant a if each agent has bounded envy ex post/ex ante with constant a of every other agent given any possible allocation. To compute this bound, we first compute all possible allocations. Then, we compute a bound that indicates how much an agent envies ex post another for each possible allocation. Now, the constant a is the greatest ex post bound over all possible allocations.

If a mechanism is envy free ex post/ex ante then it is bounded envy free ex post/ex ante with constant 0, whilst if a mechanism is (bounded) envy free ex post (with constant a) then it is (bounded) envy free ex ante (with constant a). 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 sincerely for it. We next summarize the results about fairness that we formally show in [1].

Supposing agents act sincerely, the LIKE mechanism is envy free ex ante. To see that, note that each agent bidding for an item assigns the same expected utility to all possible allocations of that particular item. The same conclusion holds also for each agent not bidding for the item. Therefore, no agent envies ex ante another one. In contrast, the LIKE mechanism is not bounded envy free ex post as an agent may be allocated much more items than another one.

EXAMPLE 3. *Suppose the fair division of m items and 2 agents. Let each agent has utility 1 for each item. There is one outcome in which agent 1 gets all the items. In this case, agent 2 assigns a utility of m units greater to the allocation of agent 1 than to their own (empty) allocation.*

As the LIKE mechanism is strategy-proof, it seems reasonable to suppose agents act sincerely. By comparison, in this setting, the BALANCED LIKE mechanism is neither envy free ex ante nor bounded envy free ex post (or even ex ante) with general utilities. Balancing the allocation of items may prevent an agent who values an item greatly from being allocated it.

EXAMPLE 4. *Consider 2 agents and 2 items, a and b. Suppose agent 1 has utility 0 for a and v for b, but agent 2 has utility 1 for a and $v-1$ for b where $v>2$. Note that both agents have the same sum of utilities. If the 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 $v-1$. As v is unbounded, agent 2 does not have bounded envy ex post or ex ante of agent 1.*

As opposed to this negative result, we show in [1] that the BALANCED LIKE mechanism is both envy free ex ante and bounded envy free ex post with constant 1, supposing 0/1 utilities and sincere agents. In here, we add on these results by discussing how does the strategic play of the agents influences fairness of our allocations with 0/1 utilities. The BALANCED LIKE mechanism is no longer envy free ex ante in this case.

EXAMPLE 5. *Suppose we are allocating items a, b and c in alphabetical order between agents 1, 2 and 3, with the following utilities. Agent 1 improves her outcome from $\frac{13}{12}$ to $\frac{9}{8}$ by bidding only for b and c.*

	a	b	c
1	1	1	1
2	1	0	1
3	1	1	0

Supposing the agents bid sincerely, the expected utility of agent 1 of any possible allocation of a is $\frac{1}{3}$. Suppose next that agent 1 does strategically not bid for a and agents 2 and 3 bid sincerely for it. Now, her expected utility of each allocation of a is $\frac{1}{2}$ compared to value of 0 which is her expected utility of her own allocation of that item. She envies them ex ante!

Note that the BALANCED LIKE mechanism is still bounded envy free ex post supposing the agents act strategically in this setting. The set of all possible allocations supposing each simple pure Nash equilibrium is a subset of the set of all possible allocations supposing sincere play. As any allocation in the latter set is envy free ex post, it follows that any allocation in the former set is also envy free ex post. Consequently, BALANCED LIKE is bounded envy free ex post and ex ante with constant 1 supposing strategic play.

To summarize, on the basis of envy-freeness, provided utilities are (not only) 0/1 (or close to these), we might consider the BALANCED LIKE mechanism to be somewhat (less) more fair than the LIKE mechanism.

5. PROPORTIONALITY

Another notion of fairness is proportionality. A mechanism is *proportional ex post* if each of the k agents receives at least $\frac{1}{k}$ of their total utility in any possible allocation. A mechanism is *proportional ex ante* if each agent receives in expectation at least $\frac{1}{k}$ of their total utility. Also, a mechanisms is *c-proportional ex post/ex ante* if there exists a constant a such that whatever the input sequence of items π, the following inequality holds

$$\frac{1}{k} \leq c \cdot \frac{u_i(\pi)}{U_i} + a \quad (1)$$

where $u_i(\pi)/U_i$ is the minimum ratio between the utility/expected utility $u_i(\pi)$ an agent i receives on π and the sum of her utilities U_i.

If a mechanism is proportional ex post/ex ante then it is c-proportional ex post/ex ante with c between 0 and 1, whilst if a mechanism is (c-proportional) proportional ex post then it is (c-proportional) proportional ex ante. It is easy to show that no randomized mechanism for indivisible items can be proportional ex post: suppose we have one indivisible item and two or more agents that bid sincerely for the item. It is possible, however, that a mechanism is proportional ex ante.

With general utilities, the LIKE mechanism is proportional ex ante supposing agents act sincerely. By comparison, the BALANCED LIKE mechanism is not proportional ex ante in general, even with just 2 agents. Balancing the allocation of items may prevent agents from being allocated items that they value greatly.

EXAMPLE 6. *Consider 2 agents and 2 items, a and b. Suppose agent 1 has utility 0 for item a and 1 for item b, but agent 2 has utility $\frac{1}{5}$ for item a and $\frac{4}{5}$ for item b. Note that the sum of the utilities for any agent is normalized to 1 unit. If agents bid sincerely, then agent 2 has an expected utility of just $\frac{1}{5}$. The latter value is strictly less than $\frac{1}{2}\cdot 1$ which is her discounted sum of utilities.*

In order to estimate the amount of this negative effect, we use c-proportionality ex ante. In fact, the BALANCED LIKE mechanism can be at worst not c-proportional ex ante for any constant c even with 2 agents. Consider the fair division setting in Example 6. Suppose that agent 2 assigns ϵ to the first item and $1-\epsilon$ to the second one. There is a unique allocation in which she gets an expected utility of ϵ, but the sum of her utilities is 1. Hence, c is unbounded as ϵ goes to zero.

Interestingly, the BALANCED LIKE mechanism becomes proportional ex ante supposing agents act sincerely and utilities are 0/1 (or close to these). This result follows from the fact that the mechanism is envy free ex ante in this case. The latter property implies proportionality.

Thus, on the basis of proportionality, we might consider the BALANCED LIKE mechanism to be somewhat as fair as (less fair than) the LIKE mechanism when utilities are (not only) 0/1 utilities (or close to these).

6. COMPETITIVENESS

A powerful technique to study online mechanisms is competitive analysis (e.g. [12]). 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, the following inequality holds

$$SW_{OPT} \leq c \cdot SW_M(\pi) + a \qquad (2)$$

where $SW_M(\pi)$ is the egalitarian/utilitarian social welfare of the mechanism M on π, and SW_{OPT} is the optimal egalitarian/utilitarian social welfare of an offline assignment that is independent on the arrival order of the items. We suppose agents bid sincerely.

The LIKE mechanism is k-competitive when the number of agents k is bounded, even with general utilities. In this setting, 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 both egalitarian and utilitarian perspectives is at worst k. For example, with 2 agents and general utilities, the LIKE mechanism is 2-competitive. That is, the expected egalitarian or utilitarian social welfare is at least 50% of the optimal (offline) allocation. In contrast, the BALANCED LIKE mechanism is not competitive with general utilities and just 2 agents. To see that, let us consider the following example.

EXAMPLE 7. *Consider the fair division of 4 items in alphabetical order between agents 1 and 2, with the following utilities and where $\epsilon > 0$ is a small positive constant. Note that the sum of the utilities for any agent is normalized to 1 unit. The optimal egalitarian (utilitarian) offline allocation gives item d to the first agent and item b to the second agent. This has an egalitarian (utilitarian) welfare of $1-\epsilon$ unit ($2-2\epsilon$ units). On the contrary,* BALANCED LIKE *achieves very small egalitarian (utilitarian) welfare of just 2ϵ (4ϵ) units.*

	a	b	c	d
1	0	ϵ	ϵ	$1-2\epsilon$
2	ϵ	$1-2\epsilon$	0	ϵ

Only when restricted to 0/1 utilities, the BALANCED LIKE mechanism is k-competitive from egalitarian perspective. Suppose k agents and k items. Let agent i bids sincerely with 1 for the first $k-i+1$ items and with 0 for the remaining ones. The egalitarian welfare is $\frac{1}{k}$ that is the expected utility of agent k whereas the optimal (offline) mechanism allocates exactly one item to each agent achieving egalitarian welfare of 1. In this binary setting, every allocation of LIKE or BALANCED LIKE achieves the optimal utilitarian welfare.

The competitive ratio supposes that the agents act sincerely. An assumption that may not always hold. For this purpose, we consider the price of anarchy, which is essentially the competitive ratio when agents bid strategically. Thus, the price of anarchy is closely related to the competitive ratio and measures how the efficiency of a decentralized system degrades due to selfish behaviour of its agents compared to imposing a centralized solution based on sincere preferences (e.g. [10]). From an egalitarian (a 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 study lower bounds of the price of anarchy of these mechanisms in [1]. For example, with 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 $\epsilon > 0$. For the BALANCED LIKE mechanism, we have the following lower bounds on the price of anarchy. With 0/1 utilities and k agents, the price of anarchy of the BALANCED LIKE mechanism from an egalitarian perspective is at least k. Similarly to the LIKE mechanism, 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 $\epsilon > 0$.

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

7. EXPERIMENTS

To determine the impact on envy-freeness, proportionality and social welfare of these mechanisms and to determine if the BALANCED LIKE mechanism outperforms the LIKE mechanism in practice, we ran some experiments. We used a wide range of problem instances: random 0/1 utilities, random Borda utilities, random general 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 [11]. For reasons of space, we are selective in our choice and report here just some of the results. These are over the classes (1) random 0/1 utilities, (2) Borda utilities from PrefLib.org and (3) random general utilities where the maximal possible utility of an agent is equal to the number of items minus one. We observed however similar trends with the other classes.

We varied the number of agents from 2 to 5 and the number of items from 2 to 10. At each data point, we sampled 100 instances, computing the optimal (offline) allocation, and all simple pure Nash equilibria by brute force. We report on three experiments. In the first one, we estimate the envy-freeness bounds. In the second one, we study the proportionality of the mechanisms. And, in the third one, we conduct empirical analysis in order to estimate their impact on the welfares. Each figure in this section contains two graphs. In the left graph of each figure, we fix the number of agents to 5 and vary the number of items. In the right graph of each figure, we fix the number of items to 10 and vary the number of agents. In both graphs of each figure, there are 4 trends corresponding to the average performance of (1) LIKE over the instances ("like"), (2) BALANCED LIKE over the instances ("balanced"), (3) BALANCED LIKE over the best simple pure Nash equilibrium of the instances ("balanced+") with respect to the egalitarian welfare and (4) BALANCED LIKE over the worst simple pure Nash equilibrium of the instances ("balanced-") with respect to the egalitarian welfare.

In Table 1, we summarize the analytical results from [1] together with the additional ones from this paper (marked with ⋆).

mechanism	LIKE		BALANCED LIKE	
utilities	binary	general	binary	general
strategy-proof	✓	✓	×	×
envy free (ex post)	×	×	×	×
envy free (ex ante)	✓	✓	✓	×
bound envy free (ex post)	×	×	✓	×
proportional (ex post)	×⋆	×⋆	×⋆	×⋆
proportional (ex ante)	✓⋆	✓⋆	✓⋆	×⋆
c-proportional (ex ante)	1⋆	1⋆	1⋆	∞⋆
c-competitive (e)	1	k	k⋆	∞
c-competitive (u)	1	k	1	∞
price of anarchy (e)	k	k	k	k
price of anarchy (u)	1	k	1	k

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

7.1 Experiment 1: Envy-freeness

Both of our mechanisms are more or less fair in some sense, depending on the utilities the agents have for the items and their strategies. Therefore, in this subsection, we study the bounds of the envy freeness ex post. These bounds also limit the envy freeness ex ante. In particular, we compute the fairness bounds of the LIKE and BALANCED LIKE mechanisms ("like" and "balanced") supposing sincere play and the BALANCED LIKE mechanism supposing the worst ("balanced-") and the best ("balanced+") simple pure Nash equilibrium . We plot arithmetic means in all graphs in this experiment. The standard deviation does not exceed 1% of its corresponding mean for each data point.

7.1.1 *Binary utilities: random*

Figure 2 reports the bounds of the fairness in the setting with random 0/1 utilities. As a very first note, the BALANCED LIKE mechanism is bounded envy free ex post with constant 1 supposing both sincere and strategic play ("balanced", "balanced-" and "balanced+"). This result is confirmed by the corresponding trends in Figure 2. On the contrary, the LIKE mechanism ("like") is increasingly unfair ex post. Its average bound increases with the number of items (left graph), starting from 2 for 2 items and reaching a bit more than 7 for 10 items. We believe the trend of this bound will increase even further provided that there are even more than 10 items available. When varying the number of agents, however, this trend is more stable (right graph), starting from nearly 6 for 2 agents, reaching almost 8 for 4 agents and decreasing smoothly till around 7 for 5 agents. Within this setting, the worst LIKE outcomes ("like") are much less fair on average than the ones of BALANCED LIKE ("balanced"). Also, the strategic play ("balanced-" or "balanced+") has a small effect on envy-freeness compared to the sincere play ("balanced").

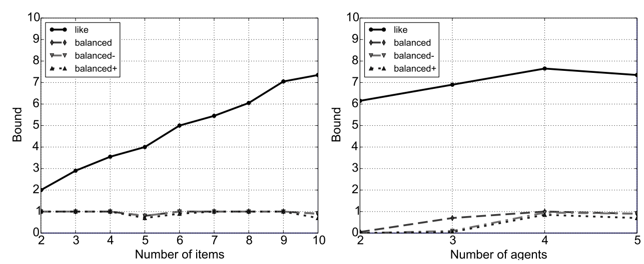

Figure 2: Envy-freeness bounds of BALANCED LIKE and LIKE mechanisms: random 0/1 utilities.

7.1.2 *Borda utilities: PrefLib*

With general utilities neither mechanism is bounded envy free ex post. Moreover, Borda utilities (e.g [2]) are a special case of general utilities in which the sum of the utilities of each agent is quadratic in the number of items. Indeed, the Borda utilities for m items are $0, 1, \ldots, m-1$ and hence their sum is equal to $m \cdot (m-1)/2$.

Figure 3 depicts two graphs. In the left graph, we can clearly see that LIKE ("like") achieves this quadratic bound for each number of items. For example, this bound is 45 for 10 items. On the contrary, the BALANCED LIKE mechanism partially restores fairness by significantly decreasing this bound. For the sake of comparison, its value is nearly 15 for 10 items (≈ 67% lower bound) supposing sincere play ("balanced") and close to 10 for the same number of items (≈ 78% lower bound) supposing strategic play ("balanced-" and "balanced+"). There is almost no difference between the bounds achieved following the best and the worst simple pure Nash equilibria in this case. In the right graph, however, we observe that LIKE mechanism ("like") achieves almost a constant bound for each number of agents that is the maximal value of 45. On the other extreme, the bound of Balanced Like mechanism starts from 0 for 2 agents and reaches slightly more than 10 for 5 agents, supposing the agents act strategically ("balanced-" and "balanced+"). In between the two extremes, it lies the fairness bound of around 15 that is achieved in case the agents are sincere ("balanced"). Strategic play seems to improve fairness with around 34% for 5 agents!

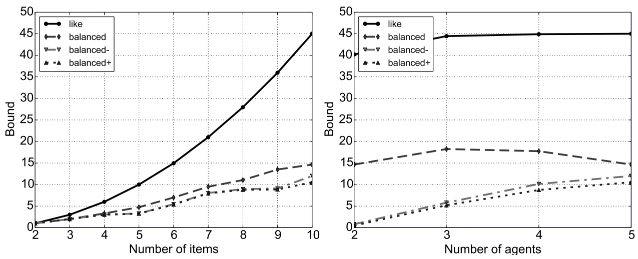

Figure 3: Envy-freeness bounds of Balanced Like and Like mechanisms: Borda utilities.

7.1.3 General utilities: random

Our mechanisms are again not bounded envy free ex post in this setting. These results are confirmed by Figure 4. The Like mechanism ("like") is less unfair in the presence of random general utilities than it was in the previous setting with Borda utilities. For instance, the maximal average bound is now around 40 for 10 items (left graph) and nearly 25 for the same number of items and 2 agents (right graph), compared to 45 achieved with Borda utilities. The trends corresponding to these bounds are however similar when we run the Balanced Like mechanism, with a noticeable improvement. The bound is 5 for 2 agents in this setting compared to 15 for the same number of agents, Borda utilities and supposing the agents are sincere ("balanced"). This is 67% less! In other words, the agents envy each other much more in the Borda setting. The fairness bounds that correspond to the strategic profiles are even lower. For 10 items, these are around 50% (≈5 against ≈10) lower in this setting than in the Borda one.

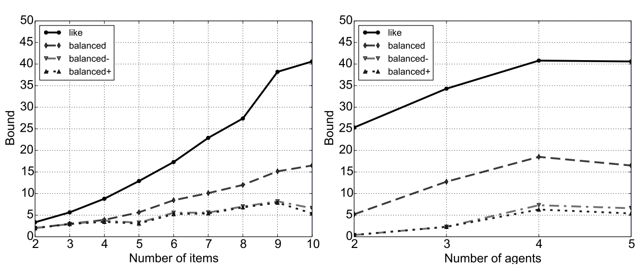

Figure 4: Envy-freeness bounds of Balanced Like and Like mechanisms: random general utilities.

7.2 Experiment 2: Proportionality

The proportionality ex ante of our mechanisms also depends on the utilities the agents have for the items. For this purpose, we focus on this property in here and compute the ratio that is reciprocal of c measuring how proportional ex ante are the Like ("like") and Balanced Like ("balanced") mechanisms over the classes of instances we select. In similar fashion, we further compute the ratios corresponding to the worst ("balanced-") and the best ("balanced+") simple pure Nash equilibria with respect to the egalitarian welfare. A value of a ratio of at least 1 indicates proportionality ex ante. We plot geometric means this time. The standard deviation is less than 1% of its respective mean for each data point.

7.2.1 Binary utilities: random

With random 0/1 utilities, both mechanisms are proportional ex ante. We note that the results in Figure 5 confirm that because the value of the ratio for each data point is greater than 1. The Balanced Like mechanism ("balanced") is nearly as proportional as Like ("like") in both graphs in Figure 5 supposing the agents are sincere. Their trends both pass over the threshold line. The value of the ratio increases once the agents are being strategic with both the number of items (left graph) and the number of agents (right graph). For example, its value is around 1.7 (1.4) for 10 items according to the best (worst) strategic play ("balanced+" ("balanced-")) compared to approximately 1.2 for the same number of items and sincere play ("balanced"). This is approximately 42% (17%) better than the sincerity. Another interesting observation at this data point is that the best ("balanced+") strategic play achieves nearly 21% more proportional allocations than the ones achieved according to the worst ("balanced-") profile.

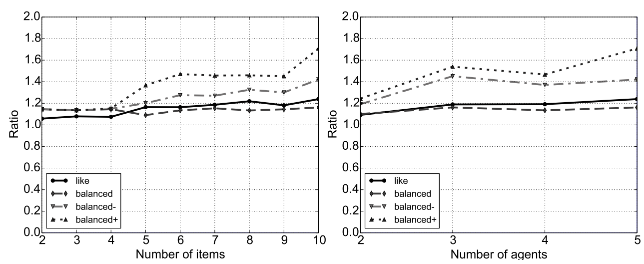

Figure 5: Proportionality of Balanced Like and Like mechanisms: random 0/1 utilities.

7.2.2 Borda utilities: PrefLib

With general utilities, Like is proportional ex ante and Balanced Like is not. There are two interesting notes in Figure 6. First, the strategic play of the agents in the best case ("balanced+") often results in as proportional allocations as the ones achieved in the worst case ("balanced-"). An exception is when the number of items is 7 (left graph). At this data point, the value of the ratio is below 1.1 in the best case compared to more than 1.2 in the worst case. Besides this possible outlier, both strategic trends improve the proportionality ratio compared to the sincere profile ("balanced"). Second, the proportionality ratio goes below the threshold of 1 supposing sincerity ("balanced") and number of items more than 7 (left graph). We believe that this trend will cross the threshold line more often in the presence of more items. When varying the number of agents (right graph), the trends corresponding to the strategic profiles ("balanced-" and "balanced+") again indicate for improved ratios compared to the sincere trend ("balanced").

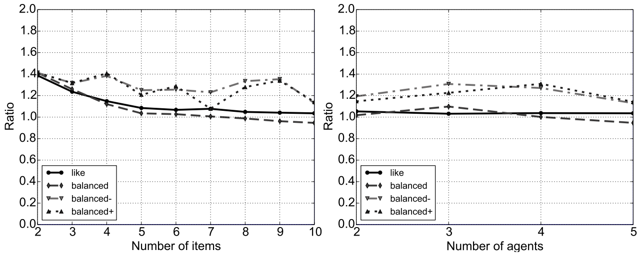

Figure 6: Proportionality of Balanced Like **and** Like **mechanisms: Borda utilities.**

7.2.3 General utilities: random

Similarly to the previous setting, Like is proportional ex ante and Balanced Like is not in this setting. Hence, one would expect that the results in here are similar to the ones in the previous setting. However, we can observe that this is not the case based on Figure 7. The strategic play of the agents improves significantly the proportionality rate this time. Both the best ("balanced+") and the worst ("balanced-") strategic profiles achieve a ratio of nearly 2 for 10 items compared to almost 1.1 that is obtained when we suppose sincere play ("balanced"). This improvement is about 82% versus only 21% achieved in the setting with Borda utilities! In addition, the trend corresponding to "balanced" passes over the unit threshold for any number of items or agents (left and right graphs). This implies that Balanced Like mechanism is proportional ex ante within this experimental setting and supposing sincere play. It competes with the Like mechanism which again is confirmed to be proportional ex ante in both graphs. To sum up, the strategic play seems to improve the proportionality of the outcomes when using Balanced Like .

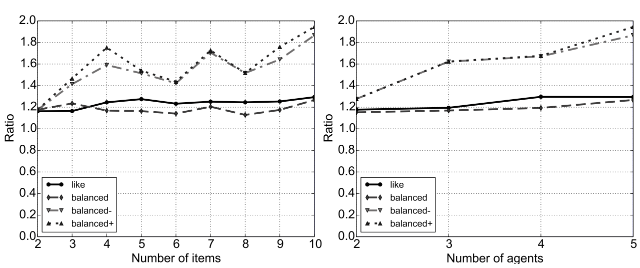

Figure 7: Proportionality of Balanced Like **and** Like **mechanisms: random general utilities.**

7.3 Experiment 3: Competitivenss

This section reports on the impact our mechanisms have on social welfare from both utilitarian and egalitarian perspective. Recall that each simple pure Nash equilibrium optimizes the utilitarian welfare with 0/1 utilities. As a response to that, we compute only the egalitarian welfare in this case. With general utilities, we compute both of them. In all settings here, we compute the reciprocals of (1) the competitive ratios ("like" and "balanced"), (2) the price of anarchy ("balanced-") and (3) the ratio between the egalitarian (utilitarian) welfare of the best simple pure Nash equilibrium and the optimal allocation ("balanced+"). As these are all ratios, we plot geometric means in all our graphs. The arithmetic means are similar. The standard deviation does not exceed 1% of its corresponding mean for each data point.

7.3.1 Binary utilities: random

Figure 8 confirms that Balanced Like ("balanced") improves the egalitarian welfare compared to Like ("like") supposing sincere or strategic play of the agents. As an example, the competitive ratio increases from around 0.7 to 0.85 on average ($\approx 21\%$ better) supposing that instead of Like ("like") we run Balanced Like ("balanced") for 10 items. Also, the strategic play of the agents often increases the welfare even in the worst case ("balanced-" compared to "balanced"), though the effect is small. With 5 agents and 10 items, the strategic profiles achieve ratios around 0.9 versus 0.85 reached by the trend of the sincere profile.

Note that each utility value of an instance from this class is 1 with probability 0.5. We believe that ranging this probability from 0 to 1 would not deliver significantly different results than the ones we report here. The rationale for that is the small welfare cost incurred when supposing strategic play.

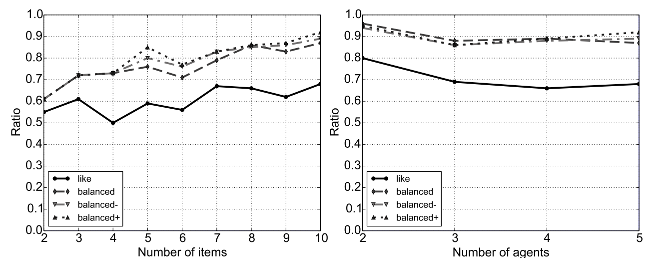

Figure 8: Egalitarian price of anarchy and competitive ratio of Balanced Like **and** Like **mechanisms.**

7.3.2 Borda utilities : PrefLib

Borda utilities are widely used in the literature (e.g. [2]) and therefore we conduct experiment in order to establish how well Like and Balanced Like perform in this setting. Figures 9 and 10 plot the respective ratios for both the egalitarian and utilitarian welfares, respectively.

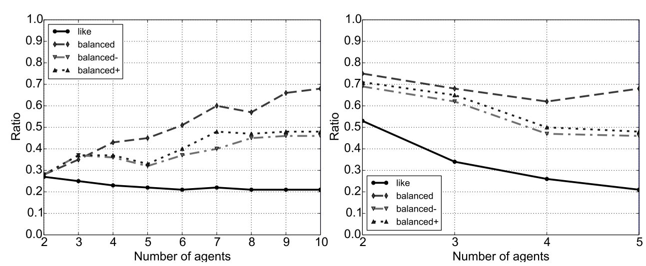

Figure 9: Egalitarian price of anarchy and competitive ratio of Balanced Like **and** Like **mechanisms.**

For the Balanced Like mechanism, one can see how the strategic play of the agents in this case ("balanced-" and "balanced+") has a strong negative impact on the egalitarian welfare compared to the case in which the agents act sincerely ("balanced"). For 10 items, the decrease is from nearly 0.7 to slightly more than 0.45. This results in approximately 36% lower egalitarian value which is a significant loss. At the same time, the worst and the best strategic profiles deliver similar egalitarian values. It seems that the agents are more strategic in the presence of general utilities than with only 0/1 utilities as now they may be more willing not to bid for items that they value low and are revealed at earlier rounds in order to bias in their favour the allocation of the following items that they may value high.

The effect the strategic play has on the utilitarian welfare is negligible in this setting as noticed in Figure 10 ("balanced-" and "balanced+" compared to "balanced"). The balanced utilitarian ratios are higher than the unbalanced ones ("like"). For 10 items, this improvement is almost 150%. To sum up, with Borda utilities, strategic play is less helpful and can result in lower egalitarian welfare. Nevertheless, LIKE remained outperformed in this setting by BALANCED LIKE .

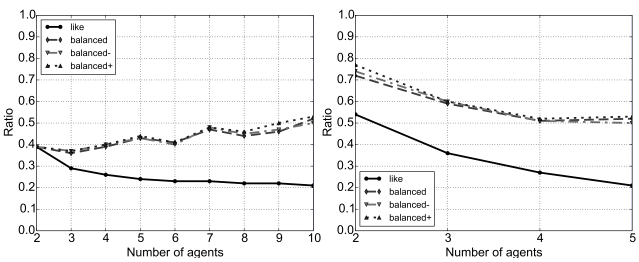

Figure 10: Utilitarian price of anarchy and competitive ratio of BALANCED LIKE and LIKE mechanisms.

7.3.3 *General utilities: random*

The last setting we present supposes the utilities are general and generated uniformly at random. In this way, the number of non-zero utility values in each instance from this class is on average the same as in an instance with Borda utilities. The difference is that now one non-zero utility value can occur multiple times whereas it occurs exactly once in a Borda profile. We believe that this class of instances is more realistic than the one with Borda utilities as normally the charities would request products they value in a similar way. This implies that the utility profiles in here would be more uniform than the Borda profiles and therefore the agents are expected to be less strategic in this setting than in the Borda setting. This intuition is confirmed by all the graphs in Figures 11 and 12.

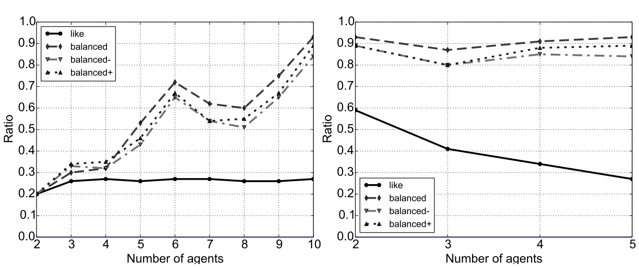

Figure 11: Egalitarian price of anarchy and competitive ratio of BALANCED LIKE and LIKE mechanisms.

In Figure 11, the ratios corresponding to the "balanced-" and "balanced+" tend to be very close to the one achieved by "balanced". The latter implies that the agents are less strategic in this setting than they were in the setting with Borda utilities. This results in smaller decrease of the egalitarian welfare in this setting than in the Borda setting. Another interesting note in this setting is that we have two clear picks in the left graph in Figure 11. These form around points corresponding to 5 and 10 items, respectively. Those values are multiple of the numbers of agents, namely 5. Due to the bias of the BALANCED LIKE mechanism, we believe that the agents get equal number of items around such points and this results in higher value of the egalitarian welfare. This effect, however, cannot be noticed in the graph on the right as there we vary the number of agents for 10 items.

We draw similar conclusions for the utilitarian welfare whose trends are shown in Figure 12. Despite the fact that this one sums up all expected utilities and, therefore, make it difficult to analyse individual allocations, we note that the utilitarian welfare is improved significantly by running BALANCED LIKE mechanism instead of LIKE mechanism. Based on both graphs, this increase is around 100% for 5 agents and 10 items.

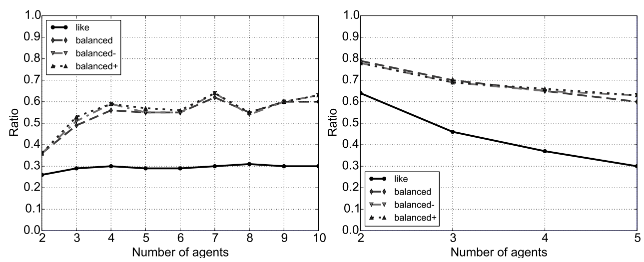

Figure 12: Utilitarian price of anarchy and competitive ratio of BALANCED LIKE and LIKE mechanisms.

8. RELATED WORK

There is a large literature on fair division of divisible and indivisible goods. Almost all studies however assume that all the goods are present initially. As opposed to that, there are a few exceptions. Walsh [14] 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 [9] have proposed a related model in which agents again arrive over time, but there are now *multiple*, *homogeneous* and *divisible* goods and not multiple, heterogeneous and indivisible goods as in here. Bounded envy freeness is closely related to the "single-unit utility difference" property that Budish, Che, Kojima and Milgrom [5] 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 (e.g [4]). These mechanisms have attracted considerable attention in the AI literature recently (e.g. [2, 8]). 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. It would be interesting future work to consider how such agent-centric mechanisms could be modified to work with online fair division problems.

9. CONCLUSIONS

Motivated by our work with local Food Bank charities, we have studied a simple online model of fair division in which a mechanism allocates packed food to agents based on their preferences. In this article, we have presented two simple such mechanisms and conducted empirical study of properties such as envy-freeness, proportionality and competitiveness. Based on our results, we might consider the BALANCED LIKE mechanism if the items can be packaged together so that agents have similar utility for all packages, and that we should otherwise prefer the LIKE mechanism when this is not possible. With similar utilities, fairness is improved by the BALANCED LIKE mechanism. This effect is strongly desired in practice as it helps quantifying both the social and individual prosperities. We established that this is the case even when the agents are strategic. In future, we plan to look into even more elaborate mechanisms than BALANCED LIKE , improving the welfares even further.

Acknowledgments

NICTA is supported by the Australian Government through the Department of Commu- nications 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)

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

[1] M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. Online Fair Division: analysing a Food Bank problem. *ArXiv e-prints*, Feb. 2015.

[2] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In T. Walsh, editor, *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pages 73–78. IJCAI/AAAI, 2011.

[3] S. Brams and A. Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, Cambridge, 1996.

[4] S. Brams and A. Taylor. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*. Norton, New York and London, 1999.

[5] 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.

[6] Y. Chevaleyre, P. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J. Padget, S. Phelps, J. Rodriguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica (Slovenia)*, 30(1):3–31, 2006.

[7] P. Davidson. *Poverty in Australia*. 2012. ISBN 978085871082.

[8] 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)*. AAAI Press, 2013.

[9] I. Kash, A. Procaccia, and N. Shah. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51:579–603, 2014.

[10] E. Koutsoupias and C. 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.

[11] 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. Springer, 2013.

[12] D. Sleator and R. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, Feb. 1985.

[13] H. Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.

[14] T. Walsh. Online cake cutting. In *Algorithmic Decision Theory, Second International Conference, ADT 2011. Proceedings*, Lecture Notes in Computer Science. Springer, 2011.

[15] 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.

[16] T. Walsh. Challenges in resource and cost allocation. In *Proceedings of the 29th AAAI Conference on AI*. Association for Advancement of AI, 2015.