

Strategic Behavior in a Decentralized Protocol for Allocating Indivisible Goods

Thomas Kalinowski, Nina Narodytska, Toby Walsh and Lirong Xia

Abstract

We study in detail a simple sequential procedure for allocating a set of indivisible goods to multiple agents. Agents take turns to pick items according to a policy. For example, in the alternating policy, agents simply alternate who picks the next item. A similar procedure has been used by Harvard Business School to allocate courses to students. We study here the impact of strategic behavior on the complete-information extensive-form game of such sequential allocation procedures. We show that computing the subgame-perfect Nash equilibrium is PSPACE-hard in general, but takes only linear time with two agents. Finally we compute the optimal policies for two agents in different settings, including when agents behave strategically and when agents can give away items.

1 Introduction

Suppose you are coaching a soccer team. To divide the players into two teams, you select the two best players as captains and then let them alternate at picking the remaining team members. Is this the best way to get an evenly matched game? Perhaps it would be better to reverse the order of their picks every round (so that the captain who picks first in the first round picks second in the second round)? This is an example of a problem in allocating indivisible goods. A number of real world problems involve allocating indivisible goods “fairly” between competing agents subject to possibly different preferences for these goods. For example, assigning courses to students at a business school is a problem of allocating indivisible goods. Students are competing for places on the popular courses, but have different preferences as to which courses to study. As a second example, the allocation of landing and take-off slots at an airport is a problem of allocating indivisible goods. Airlines are competing for popular landing and take-off times, but have different preferences as to precisely which slots they want. As a third and final example, sharing time slots on an expensive telescope is a problem of allocating indivisible goods. Astronomers are competing for observation time but have different preferences as to precisely which time slots are useful for their experiments.

Different properties might be demanded of a procedure for allocating indivisible goods. For example, we might look for allocations which are envy-free in the sense that every agent likes their allocation at least as much as the allocation to any other agent. However, envy-freeness by itself is not sufficient to ensure a “good” allocation. Not allocating any items is envy-free, and there are also many situations where no envy-free allocation exists. We might consider other criteria including efficiency of the allocation (e.g. Pareto optimality) and truthfulness of the mechanism (e.g. can agents profit by acting strategically?). There is, however, a tension between these properties. Svensson showed that the only strategy-proof, non-bossy¹ and neutral mechanism is a serial dictatorship in which agents take turns according to some order to pick their complete allocation of goods [5]. Unfortunately, a serial dictatorship can have a low efficiency in the utilitarian or egalitarian sense. In this

¹A mechanism is non-bossy if when an agent submits different preferences and their allocation does not change then the overall allocation does not change.

paper, we focus on efficiency, and consider the impact on efficiency of such issues like the strategic behavior of the agents.

2 Existing methods

Several sequential procedures for allocating indivisible goods have been proposed in the literature. For example, the Harvard Business School has been using a mechanism called *Draft* to allocate courses to students [3]. The Draft mechanism generates a priority order over all students uniformly at random. Courses are then allocated to students in rounds. In odd rounds, each student is assigned to their favourite course that still has availability using the priority order. In even rounds, the mechanism uses the reverse priority order. The Draft mechanism is not strategy-proof. Indeed, students at Harvard have been observed to behave strategically [3]. Such strategic behavior can be harmful to the ex post social welfare. However, the expected (ex ante) social welfare is higher than that of a strategy-proof mechanism like serial dictatorship. To obviate the need for certain types of manipulation, Kominers, Ruberry and Ullman [4] proposed a mechanism in which proxies play strategically. They prove that with lexicographic preferences, this proxy mechanism is Pareto efficient.

As a second example, Bouveret and Lang [1] consider a simple decentralized procedure from [2] which resembles the Draft mechanism (but ignores the initial randomisation of the order of the students). The procedure is parameterised by a policy, the sequence in which agents take turns to pick items. This policy is fixed and assumed to be known to the agents in advance. For example, as in the Draft mechanism, with two agents and four items, the policy 1221 gives first and last pick to the first agent, and second and third pick to the second agent. This procedure has the advantage that the preferences of the agents do not need to be explicitly elicited. Bouveret and Lang assume agents have additive utilities given by a common scoring function (e.g. Borda, lexicographic or quasi-indifferent scores). They consider two extreme cases: full correlation in which preference orderings of the agents are identical, and full independence in which all preference orderings are equally probable. With full correlation, all policies give the same expected sum of utilities, and the sequential allocation procedure is strategy proof. With lexicographic scores, they show that the optimal strategy for an agent given a particular policy can be computed in polynomial time supposing other agents pick truthfully. The contribution of our paper is to study this sequential allocation procedure in more detail.

3 Preliminaries

Let $\mathcal{I} = \{c_1, \dots, c_m\}$ denote a set of m indivisible goods, and $\mathcal{A} = \{A_1, \dots, A_n\}$ denote a set of n agents. For any $j \leq n$, let $u_j : \mathcal{I} \rightarrow \mathcal{R}$ denote the utility function of agent A_j over \mathcal{I} . We assume $m \geq n$, and all agents have strict preferences. That is, for any $j \leq n$ and any pair of items $\{c, c'\}$, $u_j(c) \neq u_j(c')$. We suppose that an agent's utility function is *additive*. For any $j \leq n$ and any set of items $G \subseteq \mathcal{I}$, $u_j(G) = \sum_{c \in G} u_j(c)$. For any $j \leq n$, let O_j denote the ordinal preferences of agent j . That is, O_j is a total strict order over \mathcal{I} and for any pair of items $\{c, c'\}$, $c \succ c'$ in O_j if and only if $u_j(c) > u_j(c')$. An agent has *Borda* utility, if for any $i \leq m$, the utility of the item ranked in i -th position in O_j is $m - i$. An agent has *lexicographic* utility, if for any $i \leq m$, the utility of the item ranked in i -th position in O_j is 2^{m-i} . An *allocation* is a function $f : \mathcal{I} \rightarrow \mathcal{A}$. For any agent $A \in \mathcal{A}$, $f^{-1}(A)$ denotes the set of items allocated to A . A sequential allocation is a mechanism parameterised by a *policy* P . This can be represented by an ordering over m elements taken from \mathcal{A} (e.g. $P = [A_1 \succ A_2 \succ A_1]$). Agents take turns to pick items according to this ordering.

4 Optimal Policies

Bouveret and Lang considered which policies maximise the social welfare of the agents supposing the preference of agents are independent and every preference ordering is equally likely [1]. They considered a utilitarian principle in which social welfare is measured by the expected sum of the utilities of the agents (EXPSUMUTIL). They demonstrated that the simple alternating policy 121212... optimises the social welfare when utilities are Borda score (i.e. where the i th ranked of m items has a utility of $m - i$) and up to 12 items. Interestingly, there exist situations where the policy that maximises the sum of the utilities is not alternating. In fact, it need not even be balanced (that is, it might not assign an equal number of items to both agents).

Example 1. Consider 8 items, a to h , 2 agents and utilities which are Borda scores. Suppose agent 1 has the preference order $a > \dots > h$ whilst agent 2 has the order $a > h > b > c > d > e > f > g$. Then, supposing the agents pick items truthfully, the alternating policy 12121212 gives a social welfare of $22+16=38$ but the optimal policy is 22111111 which gives a social welfare of $27+15=42$. Note that the optimal policy does not Pareto dominate the alternating policy since, whilst the optimal policy increases the utility for agent 1, the utility for agent 2 decreases slightly.

Of course, an alternating policy can still be the best policy in expectation even if there are individual situations like the above where it is not the best. Bouveret and Lang also considered an egalitarian principle in which social welfare is measured by the minimum of the expected utilities of the different agents (MINEXPUTIL). We consider two more other measures of egalitarian social welfare: the expected minimum utility of the different agents (EXPMINUTIL) and the minimum utility of the different agents over all possible worlds (MINUTIL). In the economics literature, MINEXPUTIL is called the ex-ante egalitarian utility, whilst EXPMINUTIL is called the ex-post egalitarian utility.

To illustrate the difference between the three measures, consider the following two protocols. In the first, we toss a coin. If it lands on heads, we assign all m items to agent 1, otherwise we assign all items to agent 2. In the second protocol, we assign $m/2$ items at random to agent 1 and the rest to agent 2. The second protocol is more egalitarian than the first since one agent is sure to get no items in the first protocol whilst each agent is allocated $m/2$ items in the second protocol. This is reflected in the expected minimum of the two utilities (which is zero for the first protocol and half the total utility for the second protocol), and in the minimum utility (which is zero for the first protocol, and the sum of utilities of the least valuable $m/2$ items for the second protocol). However, the minimum of the expected utilities hides this difference as both protocols have a minimum expected utility that is half the total. We have the following proposition, whose proof is straightforward and is omitted.

Proposition 1. For any policy and any distribution over utility functions:
 $\text{MINUTIL} \leq \text{EXPMINUTIL} \leq \text{MINEXPUTIL}$

Note that, whilst the minimum utility (MINUTIL) occurs in the full correlation case where agents utilities are identical [1], it can also occur when the utilities of the agents are different. For instance, suppose we are dividing just two items between two agents. Consider the protocol where the two agents declare which of the two items that they like most. If the two agents most prefer the same item, then we toss a coin to decide which agent gets this item, and assign the remaining, less preferred item to the other agent. On the other hand, if the two agents most prefers different items, we toss a coin and assign both items to an agent chosen at random. The minimum utility is now zero and occurs when the two agents most

prefer different items. The full correlation case increases MINUTIL to the smallest utility assigned to either object.

For the case of two agents, we computed the policies that maximise the three different egalitarian measures of social welfare using brute force search. Table 1 demonstrates that the optimal policies for maximising `ExpMinUtil` and `MinExpUtil` differ. We conjecture that the optimal `ExpMinUtil` policy has the form: $(12)^k 2$ for $m = 2k + 1$, $(12)^k (21)^k$ for $m = 4k$ and $(12)^k (21)^{k-1}$ for $m = 4k - 2$. In addition, we conjecture that the optimal `ExpMinUtil` policy for an even number of items is also an optimal `MinUtil` policy.

m	MINEXPUTIL	EXPMINUTIL	MINUTIL
1	1	1	1
2	12	12	12
3	122	122	122
4	1221	1221	1221
5	<i>11222</i>	12122	12122, 12212, 12211
6	121221	121221	121221, 121221, 121222, 122121, 122112
7	<i>1122122</i>	1212122	<i>1212212, 1212212, 1221122, 1221211</i>
8	<i>12212112</i>	12122121	<i>11222122, 11222211, \dots, 12212112, 12221111</i>

Table 1: Optimal policies that maximise the minimum of the two expected utilities (`MinExpUtil`), the expected minimum of the two utilities (`ExpMinUtil`) and the minimum utility (`MinUtil`). In each case, we allocate m items, assign utilities using *Borda* scoring, and assume full independence between the two agents. *Emphasis* is added to highlight when policies start to differ.

To return to our soccer example, suppose there are ten players to divide into two teams, utilities are Borda scores, and we adopt an egalitarian position to help ensure a balanced match. We might then select the two best players as team captains and, based on the optimality of the policy 12122121, have the first team captain pick first, third, sixth and eight, and the second team captain pick otherwise.

m	EXPMINUTIL egalitarian	MINUTIL egalitarian	EXPSUMUTIL utilitarian
1	1	1	1
2	12	12	12
3	122	122	121
4	1221	1222	1212
5	12122	12222	12121
6	122121	122222	121212
7	1221211	1222222	1212121
8	12212112	12222222	12121212

Table 2: Optimal policies that maximise the expected minimum of the utilities (`EXPMINUTIL`), maximise the minimum utility (`MINUTIL`) and maximise the expected sum of utilities (`EXPSUMUTIL`). In each case, we allocate m objects, assign utilities using *lexicographic* scoring, and assume full independence between the two agents.

As in [1], we also considered two other scoring models: lexicographic scoring (where an item at position k is scored 2^{-k}) and quasi-indifferent (where an item at position k is scored $a - k$ for $a \gg m$). We consider both an egalitarian model (the `EXPMINUTIL` and `MINUTIL`

m	EXPMINUTIL egalitarian	MINUTIL egalitarian	EXPSUMUTIL utilitarian
1	1	1	1
2	12	12	12
3	122	122	121
4	1221	1221	1212
5	11222	11222	12121
6	121221	121221, 122112, 122121	121212
7	1112222	1112222	1212121
8	12122121	11222211, 12122121, 12211221, 1221211	12121212

Table 3: Optimal policies that maximise the expected minimum of the utilities (EXPMINUTIL), maximise the minimum utility (MINUTIL) and maximise the expected sum of utilities (EXPSUMUTIL). In each case, we allocate m objects, assign utilities using quasi-indifferent scoring, and assume full independence between the two agents.

policies in which we maximise the expected or actual minimum utilities) and a utilitarian model (the EXPSUMUTIL policy in which we maximise the expected sum of the utilities). In Tables 2 and 3, we report the optimal policies for lexicographic and quasi-indifferent scoring.

We make some observations about these results. First, in both scoring models, a simple alternating policy is optimal under the utilitarian assumption. It seems likely that the expected sum of utilities is maximised for a wide variety of scoring functions by this policy. Second, for the quasi-indifferent scoring function, the same policy is optimal for EXPMINUTIL and MINUTIL. This was not the case for the lexicographic scoring model. For Borda scoring, the same policy was optimal for EXPMINUTIL and MINUTIL only for even n .

5 Strategic Behavior

Another desirable property of an allocation procedure is strategy-proofness. A sequential allocation procedure is strategy-proof if for any utility functions, the agents are best off choosing their top ranked item still available at every step. Unfortunately, the sequential allocation procedure is not strategy-proof in general. For instance, the first agent to pick an item might not pick their most preferred item if this is the item least preferred by the other agent. The first agent might strategically pick some other item as the second agent will not pick this first item unless there is no other choice. Bouveret and Lang [1] argue that the sequential allocation procedure is strategy-proof when agents have the same preference rankings. They also gave a polynomial time method for a single agent to compute a manipulation supposing all other agents act truthfully and utilities are lexicographic. Supposing all agents but the manipulator act truthfully is a strong assumption. If one agent is acting strategically, why not the others?

The sequential allocation procedure naturally lends itself to a game theoretic analysis in which all agents can act strategically. Assuming that the agents know the utility functions of other agents, we can model the sequential allocation procedure as a complete information *extensive-form game*. The subgame-perfect Nash equilibrium (SPNE) gives the (perhaps untruthful) strategy in which agents cannot improve their allocation by deviating unilaterally. The SPNE can be computed by *backward-induction* as follows. We start with the last agent A in the order P . For any allocation of items in the previous rounds, only one item remains, and A will get it. Then, we move to the second to the last agent A' in P .

For any allocation of items in previous round, A' can predict the final allocation for any item she picks. Therefore, she can pick an item that maximises her total utility in the final allocation. We then move on to the third to the last agent in P , etc. Since an agent can obtain the same total utility for picking different items, there might be multiple SPNE.

Example 2. *Suppose there are two agents and four items. Agent 1's ordinal preferences are $O_1 = c_1 \succ c_2 \succ c_3 \succ c_4$ and agent 2's ordinal preferences are $O_2 = c_2 \succ c_3 \succ c_4 \succ c_1$. Let $P = A_1 \succ A_2 \succ A_2 \succ A_1$. If all agents behave truthfully, then A_1 chooses c_1 in the first round, A_2 chooses c_2 and c_3 in the second and third rounds, respectively, and A_1 chooses c_4 in the last round. If the agents behave strategically, then A_1 can choose c_2 in the first round, and still get c_1 in the last round. The unique SPNE allocation in this game has A_1 getting $\{c_1, c_2\}$ and A_2 getting $\{c_3, c_4\}$.*

In the above example, even though there are multiple SPNE, the final allocation is unique regardless of the utility functions. We will see later that this is not a coincidence. When there are two agents, the *final* SPNE allocation is always unique (and indeed can be computed in linear time). The next example shows that with three or more agents, there can be multiple SPNE allocations.

Example 3. *Suppose there are four items and three agents with Borda utilities. The ordinal preferences of the agents are as follows. $A_1 : c_1 \succ c_2 \succ c_3 \succ c_4$, $A_2 : c_3 \succ c_4 \succ \dots$, and $A_3 : c_1 \succ c_2 \succ \dots$. Let $P = A_1 \succ A_2 \succ A_3 \succ A_1$. There are two SPNE allocations: (1) if A_1 picks c_1 in the first round, then in the SPNE A_1 gets $\{c_1, c_4\}$, A_2 gets c_3 , and A_3 gets c_2 ; (2) if A_1 picks c_3 in the first round, then in the SPNE A_1 gets $\{c_2, c_3\}$, A_2 gets c_4 , and A_3 gets c_1 .*

5.1 Computing SPNE for Two Agents

With two agents and m items, computing the subgame-perfect Nash equilibrium by backward induction takes $\Omega(m!)$ time. This will be prohibitive when we have many items. The SPNE can, however, be computed in just $O(m)$ time by means of the following result. Let u_1, u_2 be the utility functions of the two agents, O_1, O_2 be their ordinal preferences, and P be the policy. We let $\text{Seq}(O_1, O_2, P)$ denote the truthful sequential allocation. We use $\text{SPNE}(u_1, u_2, P)$ to denote the subgame-perfect Nash equilibrium allocation. For any total strict order O , let $\text{rev}(O)$ denote the reversed order. Then, we can show that the SPNE allocation is unique, and can be computed from the truthful sequential allocation for the reversed preference orderings and policy.

Theorem 1. *When there are two agents, the SPNE allocation is unique. Moreover, $\text{SPNE}(u_1, u_2, P) = \text{Seq}(\text{rev}(O_2), \text{rev}(O_1), \text{rev}(P))$*

Proof: (Sketch) W.l.o.g. suppose agent 1 has the last pick in policy P (and thus the first pick in policy $\text{rev}(P)$). Then, agent 1 knows that the item that is ranked last in O_2 is “safe”, as agent 2 has no incentive to pick it in earlier rounds. Therefore, agent 1 can safely pick this item in her last round, and leave opportunities in previous rounds in P to pick more popular items. The formal proof is much more involved and is proved by induction on the number of items m . ♣

Example 4. *Suppose there are two agents and four items. The agents' preferences and the policy are the same as in Example 2. We have $\text{rev}(P) = P$. In $\text{Seq}(\text{rev}(O_2), \text{rev}(O_1), \text{rev}(P))$, A_1 picks c_1 in the first round, A_2 picks c_3 and c_4 in the second round and third round respectively, and A_1 picks c_2 in the last round. This outcome is the same as the SPNE allocation in Example 2.*

5.2 Computing SPNE for more than Two Agents

When the number of agents n is comparable to the number of items m (more precisely, when $n = O(m)$), we prove that computing the SPNE is intractable. Consider the decision problem SUBGAMEPERFECT, where we are given the utility functions of n agents over m items, a particular agent A , a policy P , and a threshold T , and we are asked whether the utility of A is larger than T in any SPNE. Computing the SPNE for a finite multi-player extensive game with perfect information is naturally in PSPACE [6]. Our contribution here is to show that this particular game is complete for this complexity class.

Theorem 2. SUBGAMEPERFECT is PSPACE-complete for Borda scoring of utilities.

Proof: Backward induction shows that it is in PSPACE. To show hardness, we give a reduction from QSAT, which is a standard PSPACE-complete problem. In a QSAT instance, we are given a quantified formula $\exists x_1 \forall x_2 \exists x_3 \cdots \forall x_q . \varphi$ where q is even and we are asked whether the formula is true. Let $\varphi = C^1 \wedge \cdots \wedge C^t$, where C^j is a 3-clause, $l_j^1 \vee l_j^2 \vee l_j^3$. We construct a SUBGAMEPERFECT instance where there is a unique SPNE with a utility to the first player larger than a threshold if and only if the formula is true.

In the SUBGAMEPERFECT instance, there are q agents who represent the binary variables. Each of these agents choosing one out of two items represents a valuation of the variable. The agents that correspond to \exists quantifiers (that is, agents 1, 3, ..., $q-1$) obtain higher utility if φ is true under the current valuation, and the agents that correspond to \forall quantifiers (that is, agents 2, 4, ..., q) obtain higher utility if φ is false under the current valuation. There are also some other agents that are used to encode the QSAT instance, which we will specify later.

Let a be an item, and k, p be natural numbers. We define an ordering $O_p^k(a)$ that will be used as part of the policy P as follows. It introduces $2k+1$ new agents A_p^1, \dots, A_p^{2k+1} and $5k+1$ new items $\{a_p, b_p^1, \dots, b_p^k, c_p^1, \dots, c_p^k, d_p^1, \dots, d_p^k, e_p^1, \dots, e_p^k, f_p^1, \dots, f_p^k\}$. The preferences of the new agents are as follows:

Agent	Preferences
A_p^1	$b_p^1 \succ c_p^1 \succ d_p^1 \succ e_p^1 \succ \text{Others}$
\vdots	\vdots
A_p^k	$b_p^k \succ c_p^k \succ d_p^k \succ e_p^k \succ \text{Others}$
A_p^{k+1}	$c_p^1 \succ f_p^1 \succ \text{Others}$
\vdots	\vdots
A_p^{2k}	$c_p^k \succ f_p^k \succ \text{Others}$
A_p^{2k+1}	$a \succ b_p^k \succ \cdots \succ b_p^1 \succ a_p \succ \text{Others}$

Let the order over agents be $A_p^1 \succ \cdots \succ A_p^{2k+1} \succ A_p^1 \succ \cdots \succ A_p^{2k}$. In $O_q^k(a)$, a is the item that we want to “duplicate”, k is the number of duplicates, and q is merely an index. We can prove by induction that if a has not been chosen (in previous rounds), then after agents have chosen items according to $O_q^k(a)$, $\{f_p^1, \dots, f_p^k\}$ will be chosen and $\{d_p^1, \dots, d_p^k\}$ will not be chosen; if a has been chosen, then $\{d_p^1, \dots, d_p^k\}$ will be chosen rather than $\{f_p^1, \dots, f_p^k\}$.

We now specify the sequential allocation instance by using the orderings $O_p^k(a)$. All agents introduced in $O_p^k(a)$ will not appear in other places in the policy P . For each $i \leq q$, there are two items 0_i and 1_i that represent the two values of x_i , an agent A_i corresponding to the valuation and another agent B_i that is used to make sure that A_i chooses 0_i or 1_i in the $(q+2i-1)$ th round. For each $i \leq q$, D_i is an agent whose preferences are $d_i \succ \text{Others}$, where d_i is a new item that creates a “gap” between items available to agent A_i . The first $(2t+4)q$ agents in P are the following: $D_1 \succ \cdots \succ D_q \succ A_1 \succ \cdots \succ A_q \succ O_1^t(0_1) \succ \cdots \succ$

$O_q^t(0_q) \succ B_1 \succ \dots \succ B_q$. The preferences of B_i are $0_i \succ 1_i \succ \text{Others}$. The preferences of A_i will be defined after we have defined all items and have specified P . For notational convenience, for each $i \leq q$ and each $j \leq t$ we rename d_i^j to be 0_i^j , and rename f_i^j to be 1_i^j .

For each clause C^i , we have an agent denoted by C_i . Suppose v_{j_1}, v_{j_2} , and v_{j_3} correspond to the 3 valuations that satisfy C_i , then we let the preferences of C_i be $v_{j_1}^i \succ v_{j_2}^i \succ v_{j_3}^i \succ g \succ g'_i \succ \text{Others}$, where g and g'_i are new items. g is used to detect whether a clause is not satisfied. For example, suppose $C^i = x_1 \vee \neg x_2 \vee x_3$, then the preferences of C_i are $1_1^i \succ 0_2^i \succ 1_3^i \succ g \succ g'_i \succ \text{Others}$. The remaining agents in the P are: $C_1 \succ \dots \succ C_t \succ O_{q+1}^q(g) \succ A_1 \succ \dots \succ A_q$.

The agents and new items introduced in $O_{q+1}^q(g)$ impose “feedback” on A_1 through A_q , such that if g is allocated before $O_{q+1}^q(g)$ (which means that the formula is not satisfied under the valuation encoded in the first q rounds), then some items that are more valuable to the agents that correspond to the \forall quantifiers are made available; if g is not allocated before $O_{q+1}^q(g)$, then some items that are more valuable to the agents that correspond to the \exists quantifiers are made available. Finally, for each $i \leq q$, we define the ordinal preferences of A_i as follows. If i is odd, then A_i 's preferences are $0_i \succ 1_i \succ d_{q+1}^i \succ d_i \succ f_{q+1}^i \succ \dots$. If i is even, then A_i 's preferences are $0_i \succ 1_i \succ f_{q+1}^i \succ d_i \succ d_{q+1}^i \succ \dots$.

To summarise, in the sequential allocation instance, there are $3q + t + (2t + 1)q + 2q + 1$ agents and $m = 3q + (5t + 1)q + 1 + t + 5q + 1$ items, which are polynomial in the size of the formula ($\Omega(t + q)$). Table 4 summaries the items introduced in the reduction. Final, the

for	items	Introduced in
$i \leq q$	d_i	D_i
$i \leq q$	$0_i, 1_i$	A_i
$i \leq q, j \leq t$	a_i b_i^j c_i^j d_i^j (a.k.a. 0_i^j) e_i^j f_i^j (a.k.a. 1_i^j)	$O_i^j(0_i)$
	g	C_1
$j \leq t$	g'_t	C_j
$j \leq q$	$a_{q+1}, b_{q+1}^j, c_{q+1}^j,$ $d_{q+1}^j, e_{q+1}^j, f_{q+1}^j$	$O_{q+1}^q(g)$

Table 4: Items introduced in the reduction.

policy P ordering over agents is the following.

$$\begin{aligned}
D_1 \succ \dots \succ D_q \succ A_1 \succ \dots \succ A_q \succ O_1^t(0_1) \succ \dots \succ O_q^t(0_q) \\
\succ B_1 \succ \dots \succ B_q \succ C_1 \succ \dots \succ C_t \succ O_{q+1}^q(g) \\
\succ A_1 \succ \dots \succ A_q
\end{aligned}$$

If we must allocate all items then we can add some dummy agents to the end of the ordering.

We note that if an agent only appears once in the ordering, then it is her strictly dominant strategy to pick her most preferred available item. In any SPNE, in the first q rounds d_1, \dots, d_q will be chosen. In the next q rounds, agent i must choose either 0_i or 1_i , otherwise 0_i will be chosen by agent A_i^{2t+1} introduced in $O_i^t(0_i)$ and 1_i will be chosen by B_i . Hence, the choices of agents A_i correspond to valuations of the variables, and these valuations are duplicated by $O_i^t(0_i)$ that will be used to satisfy clauses. (We note that if A_i chooses 0_i ,

then after $O_i^t(0_i)$, $\{0_i^1, \dots, 0_i^t\}$ are still available, but $\{1_i^1, \dots, 1_i^t\}$ are not available; and vice versa.) Then, a clause C^i is satisfied if and only if at least one of the top 3 items of agent C_i is available (otherwise C_i chooses g). Hence, after agent C_t , g is available if and only if all clauses are satisfied. Finally, if g is available after agent C_t , then the agents that correspond to the \exists quantifiers can choose d_{q+1} 's to increase their total utility by $m - 3$, but the agents that correspond to the \forall quantifiers can only choose d_{q+1} 's to increase their utility by $m - 5$; and vice versa. Hence, the agents that correspond to \exists quantifiers will choose valuations to make F true, while the agents that correspond to \forall quantifiers will choose valuations to make F false. It can be verified that there is a unique SPNE allocation, where agent A_1 's utility is at least $2m - 5$ (that is, she gets one of $\{0_1, 1_1\}$ and d_{q+1}^1) if and only if the formula F is true. ♣

6 Optimal Policies for Strategic Behavior

Suppose agents act strategically instead of truthfully. For example, suppose they pick items according to the subgame-perfect Nash equilibrium. The policies which maximise social welfare can now change. For a reversal symmetric scoring function like Borda, and a reversal symmetric policy like the simple alternating policy, the situations where strategic behavior decreases social welfare are exactly balanced by the symmetric situations where it increases social welfare. As a result, we did not observe any difference in the policies that optimises social welfare for Borda scoring when agents behave strategically instead of truthfully. For example, brute force calculation with up to 8 items show that the expected sum of the utilities of the agents supposing Borda scoring is maximised by the same simple alternating policy whether agents pick either truthfully or strategically.

Strategic behavior can sometimes increase the social welfare of the agents. In other cases, it can decrease the social welfare of the agents or leave it unchanged. In fact, given the reversal symmetry of the optimal policy, and of the subgame perfect equilibrium, Borda scoring and the utilitarian criterion, we can prove that the cases when the utilitarian social welfare increases are exactly matched by cases where it decreases. With an egalitarian criterion, strategic behavior can improve social welfare slightly more often than it can decrease it. Averaged over all possible preference profiles, brute force calculations suggest that the expected sum of the utilities barely changes, whilst the expected minimum increases by less than 1%.

For scoring functions that are not symmetric, the optimal policy can change. For example, with lexicographic scores, the optimal policy for strategic behavior is different from that for truthful behavior. Table 5 summarises results based on brute force calculation. When maximising the expected minimum utility, the optimal policies for agents playing strategically are optimal policies for agents playing truthfully for 6 or fewer items. However, the optimal policy for strategic play with 7 items is 1221122 but for truthful play is 1221211. Similarly, for 8 items, the optimal policy for strategic play is 12212211 but for truthful play is 12212112. When maximising the minimum utility, the optimal policies for strategic play are optimal policies for truthful play. When maximising the expected sum of utilities and 4 or more items, the optimal policies for strategic play are not optimal alternating policies for truthful play.

We conjecture that the optimal `ExpMinUtil` policy supposing strategic behavior has the alternating form: $(1221)^k 21$ for $m = 4k + 2$, $(1221)^k 122$ for $m = 4k + 3$ and $(1221)^k 2211$ for $m = 4k + 4$. We also conjecture that the optimal `ExpSumUtil` policy supposing strategic behavior has the alternating form: $(12)^k 122$ for $m = 2k + 3$, $1(2211)^k 2$ for $m = 4k + 2$, and $1(2211)^k 221$ for $m = 4k + 4$. Strategic play also carries a small cost. Averaged over all possible preference profiles, the utility decreases by 5% or less for both the expected sum

n	ExpMinUtil egalitarian	MinUtil egalitarian	ExpSumUtil utilitarian
1	1	1	1
2	12	12	12
3	122	122	121
4	1221	1222	1212, <i>1221</i>
5	12122	12222	<i>12122</i>
6	122121	122222	<i>122112</i>
7	<i>1221122</i>	1222222	<i>1212122</i>
8	<i>12212211</i>	12222222	<i>12211221</i>

Table 5: Optimal policies when we assign utilities using lexicographic scoring, and assume agents play strategically by computing the subgame-perfect Nash equilibrium. *Emphasis* is added to highlight when policies differ from the optimal truthful policies.

and minimum of utilities.

As in [1], we also considered quasi-indifferent scoring. With quasi-indifferent scoring, an item at position k in an agent's ordering is given score $a - k$ where $a \gg m$. In Table 6, we give the optimal policies for agents playing strategically when agents are quasi-indifferent between items. The optimal policy for agents playing strategically is also the optimal policy for agents playing truthfully except $m = 6$ and the egalitarian criterion of maximising the expected minimum utility. When agents play strategically, the optimal policy in this case is 122121. However, when agents play truthfully, the optimal policy in this case is 121221.

m	ExpMinUtil egalitarian	MinUtil egalitarian	ExpSumUtil utilitarian
1	1	1	1
2	12	12	12
3	122	122	121
4	1221	1221	1212
5	11222	11222	12121
6	122121	121221, 122112, 122121	121212
7	1112222	1112222	1212121
8	12122121	12211221, 12212112, 12122121, 11222211	12121212

Table 6: Optimal policies when we assign utilities using a quasi-indifferent scoring function, and assume agents play strategically by computing the subgame perfect equilibrium.

7 Disposal of Items

One inefficiency of the policies considered so far is that one agent may use one of their early choices to select an item that the other agent would happily give away. There is an inherent asymmetry in agents declaring items that they like most but not the items that they like least. To address this issue, we suppose agents can select the item that they least like to give to the other agent. For instance, the policy $1\bar{1}21$ describes a protocol in which the first agent starts by picking their most preferred item, then picks their least preferred item to give to the second agent, the second agent then picks the most preferred of the two items that remain, and the first agent then gets the last remaining item. $\bar{1}$ means that agent

1 gives the item remaining that she likes least to agent 2. Such disposal of items can be extended to more than 2 agents but requires a protocol for which agent takes the disposed item.

m	ExpMinUtil egalitarian	ExpSumUtil utilitarian
1	1	1
2	12	12
3	122	121, $\bar{1}21$
4	1221, $1\bar{1}21$, $1\bar{2}22$, $\bar{1}211$	$1\bar{1}21$
5	12122, $1\bar{1}\bar{2}12$	12 $\bar{2}12$, $1\bar{2}\bar{2}12$
6	$1\bar{2}\bar{1}\bar{1}21$, $\bar{1}21121$	$1\bar{1}2121$, $1\bar{1}\bar{2}\bar{1}21$
7	12 $\bar{1}\bar{1}\bar{2}12$	1212 $\bar{2}12$, $1\bar{2}\bar{1}2\bar{2}12$, $\bar{1}212\bar{2}12$, $1\bar{2}\bar{1}2\bar{2}12$
8	$1\bar{2}\bar{1}\bar{1}\bar{2}\bar{1}21$, $\bar{1}2112121$	$1\bar{1}212121$, $1\bar{1}21\bar{2}\bar{1}21$, $1\bar{1}\bar{2}\bar{1}2121$, $1\bar{1}\bar{2}\bar{1}\bar{2}\bar{1}21$

Table 7: Optimal policy for dividing m items with utility measured using Borda scoring assuming egalitarianism or utilitarianism and full independence between the two agents. Note that when computing the optimal policy, we consider all possible policies including those in which agents only pick items, and those in which agents only give items away.

In Table 7, we give the optimal policies assuming strategic behavior, and Borda scoring of utilities when agents can dispose of items as well as pick them. We again put policies into a canonical form in which agent 1 makes the first move. There is a symmetric policy in which we swap agent 1 with agent 2 throughout. We also ignore policies which result in the same division of items. For instance, a policy containing the moves $\bar{1}1$ is equivalent one containing $1\bar{1}$. Our canonical form has agents picking items before giving give them away. For example, a policy that ends with the moves $\bar{2}1$ gives the last two items to the first agent so is equivalent to one that ends with the moves 11 . Our canonical form describes a policy by the lexicographically least equivalent policy supposing that 1 and 2 are ordered before $\bar{1}$ and $\bar{2}$.

We make some observations about the results. First, we can often increase social welfare by having agents declare items that they dislike. There are a few optimal policies in which agents only pick items that they like (e.g. for $m = 5$, one of the optimal egalitarian policies is 12122). However, in most cases, the optimal policy has agents declaring both items that they like and dislike. Second, when dividing 4 items between two agents, there is a policy, $1\bar{1}21$ that is optimal for both the egalitarian and utilitarian measures of social welfare. Third, unlike protocols in which agents pick just items that they like, there are often several different protocols which maximise social welfare.

8 Conclusions

We have studied a simple sequential allocation procedure where agents get to choose items according to a policy, and agents have simple additive utilities over items given by Borda, lexicographic or quasi-indifferent scores. We have computed optimal policies assuming both truthful and strategic behavior of the agents for both egalitarian and utilitarian measure of social welfare. We have also proved that with two agents, the subgame perfect Nash equilibrium is polynomial to compute by simply reversing the agents' preferences and the policy. On the other hand, with more than two agents, we proved that computing the subgame perfect Nash equilibrium is PSPACE-hard. There are many directions for future work. One direction would be to prove the conjectures about the optimal policies for maximising social welfare assuming truthful or strategic behavior and Borda or lexicographic scoring. Another

direction would be to determine if we can compute the subgame-perfect Nash equilibrium in polynomial time for a fixed number agents k where $k > 2$. More generally, when we want to allocate multiple indivisible goods, how can we design simple decentralized mechanisms that balance efficiency and strategy-proofness?

References

- [1] 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.
- [2] S.J. Brams and A.D. Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, Cambridge, 1996.
- [3] E. Budish and E. Cantillon. The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard. *American Economic Review*, page (forthcoming), 2012.
- [4] S.D. Kominers, M. Ruberry, and J. Ullman. Course allocation by proxy auction. In A. Saberi, editor, *6th International Workshop on Internet and Network Economics (WINE 2010)*, volume 6484 of *Lecture Notes in Computer Science*, pages 551–558. Springer, 2010.
- [5] L.-G. Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.
- [6] C. Ivarez, J. Gabarr, and M. Serna. Polynomial space suffices for deciding Nash equilibria properties for extensive games with large trees,. In X. Deng and D.-Z. Du, editors, *Algorithms and Computation*, volume 3827 of *Lecture Notes in Computer Science*, pages 634–643. 2005.

Thomas Kalinowski
Universität Rostock
Rostock, Germany
Email: `thomas.kalinowski@uni-rostock.de`

Nina Nardodytska and Toby Walsh
NICTA and UNSW
Sydney, Australia
Email: `{nina.narodytska,toby.walsh}@nicta.com.au`

Lirong Xia
Harvard University
Cambridge, USA
Email: `xialirong@gmail.com`