

Welfare of Sequential Allocation Mechanisms for Indivisible Goods

Haris Aziz¹ and Thomas Kalinowski² and Toby Walsh³ and Lirong Xia⁴

Abstract. Sequential allocation is a simple and attractive mechanism for the allocation of indivisible goods used in a number of real world settings. In sequential allocation, agents pick items according to a policy, the order in which agents take turns. Sequential allocation will return an allocation which is Pareto efficient – no agent can do better without others doing worse. However, sequential allocation may not return the outcome that optimizes the social welfare. We consider therefore the relationship between the welfare and the efficiency of the allocations returned by sequential allocation mechanisms. We then study some simple computational questions about what welfare is possible or necessary depending on the choice of policy. Over half the problems we study turn out to be tractable, and we give polynomial time algorithms to compute them. We also consider a novel control problem in which the Chair chooses a policy to improve social welfare. Again, many of the control problems we study turn out to be tractable, and our results give polynomial time algorithms. In this case, tractability is a good thing so that the Chair can improve the social welfare of the allocation.

INTRODUCTION

The fair division of resources is a central problem in social choice. One challenging case in fair division is when the goods being allocated are indivisible. For instance, we might want be interested in allocating courses to students at an university, time windows on an expensive scientific instrument to different groups of scientists, or landing slots on a runway to different airlines.

A simple mechanism to allocate indivisible goods like this is *sequential allocation* [9]: agents simply take turns to pick items. This leaves open the particular order (the “policy”) used to pick items. For example, in a balanced alternating policy, agents pick items in rounds, every agent picks one item in each round, and the order of agents is reversed between rounds. On the other hand, in a balanced policy, the agents simply have the same number of turns, and there is no restriction that picking happens in rounds with each agent getting one item in each round. Throughout this paper to make life simple and ensure balance is indeed possible, we will assume that the number of items is an integer multiple of the number of agents.

The actual policy used may not be fixed in advance. For example, sequential allocation is used to allocate courses to students at the Harvard Business School [10], and the policy used is chosen uniformly from the space of all balanced alternating policies by randomly ordering the students in the initial round. Whilst this may be perceived

to be procedurally fair, it does not necessarily maximize the welfare of the agents.

This suggests a number of questions about the social welfare that can or must be achieved by sequential allocation mechanisms. Do we *necessarily* achieve a minimum acceptable welfare whatever policy is chosen? Is it *possible* that the welfare is above some given amount? What is the *maximum* or *minimum* welfare that can be achieved? These questions are closely related to an interesting control problem. Can a (benevolent) chair choose a policy not at random but to improve or maximize welfare? They are also related to the expected welfare when the policy is chosen at random, as at the Harvard Business School. The expected welfare is between the minimum welfare that is necessary and the maximum welfare that is possible. Indeed, if the minimum and maximum welfare are different, then the expected welfare is strictly between them.

We study these problems about the welfare possible or necessary from a computational perspective. We consider classes of policies considered in previous work (e.g. [8, 5, 1]), and used in real life setting like the previously mentioned course allocation mechanism from Harvard Business School. As our results show (summarized in Table 1 at the end of the paper), over half of these problems are polynomial time solvable. This is a good thing. We want to be able to compute the welfare possible or necessary. We want the Chair potentially to be able to improve the welfare by choosing a good policy. We want to be able to compute the expected welfare of the agents. Our results provide efficient algorithms to compute answers to these questions.

Sequential allocation is an ordinal mechanism. That is, it merely requires agents to declare an ordering over items. It does not require the agents to declare their actual utilities. However, to compute the welfare of an outcome, we need to know the utilities of the agents. This does not necessarily mean we need to elicit the utilities explicitly. A significant body of work in the fair division literature supposes agents have utilities that are simply derived from their ordinal preferences (e.g. an agent’s utility is simply the sum of the Borda scores for the items). For example, Brams *et al.* [7] study fairness criteria using, amongst others, Borda utilities derived from the ordinal preferences. As a second example, Bouveret and Lang [5] consider sequential allocation mechanisms with Borda, lexicographical or quasi-indifferent utilities. As a third example, Kalinowski *et al.* [13] compute the optimal policies for sequential allocation supposing agents have Borda utilities. As a fourth example, Baumeister *et al.* [3] study ordinal mechanisms based on Borda, lexicographical, quasi-indifferent or *k*-approval utilities. As a fifth example, Darmann and Schauer [11] consider mechanisms that maximize the Nash product social welfare supposing Borda, lexicographical or 0/1 utilities. As a sixth example, Fujita *et al.* [12] study mechanisms for house allocation supposing lexicographical utilities. Many of our results hold

¹ Data61 and UNSW, email: haris.aziz@data61.csiro.au

² University of Newcastle, email: thomaskalinow@gmail.com

³ UNSW and Data61, email: toby.walsh@data61.csiro.au

⁴ RPI, email: xialirong@gmail.com

Toby Walsh is now funded by the European Research Council under the EU's Horizon 2020 programme via AMPLIFY 670077.

in the special case of Borda or lexicographical utilities, and thus only require agents to declare ordinal preferences.

Even if we elicit general utilities as opposed to ordinal preferences, there are some advantages to an ordinal mechanism like sequential allocation. First, it is easier for agents to verify that the sequential allocation mechanism has been applied correctly compared to, say, a black-box cardinal mechanism that maximizes welfare. The chair can nevertheless still choose a policy that maximizes welfare. Second, the sequential allocation mechanism can easily ensure additional constraints like, for instance, that all agents receive the same number of items, or that an agent is not allocated two items that are incompatible with each other. Such constraints can significantly increase the complexity of applying a cardinal mechanism. Third, sequential allocation mechanisms, especially when restricted to a class like balanced alternation policies, may be perceived to be procedurally fair. Fourth, there is typically less opportunity for agents to act strategically with an ordinal mechanism like sequential allocation.

To demonstrate that ordinal mechanisms may offer less opportunity for strategic behaviour than cardinal mechanisms, consider 2 agents and 4 items, a to d . Suppose agent 1 sincerely declares Borda utilities: 4 for a , 3 for b , 2 for c and 1 for d . If we use the sequential allocation procedure with a balanced alternation policy then agent 1 always gets one of its top two choices or both, irrespective of how agent 2 acts, strategically or sincerely. Consider now the cardinal mechanism that maximizes the utilitarian welfare. Suppose agent 2 has the same Borda utilities as agent 1 but strategically declares an utility of 5 for a , 4 for b , 1 for c and 0 for d . Note that the total utility declared by agent 2 is the same as the sum of the Borda scores of agent 1. With these declared utilities and the cardinal mechanism that maximizes the utilitarian welfare, agent 2 now gets both of agent 1's top two choices. Thus, we see more strategic outcomes are possible with this cardinal mechanism.

WELFARE AND EFFICIENCY

When agents pick sincerely, sequential allocation is guaranteed to return a Pareto efficient outcome. No agent can do better without at least one being worse off. However, sequential allocation is not guaranteed to maximize the social welfare of the outcome. We consider therefore the precise relationship between social welfare and efficiency. We suppose that there are n agents being allocated $m = nk$ items for integer $k \geq 1$. Agents have additive utilities over the items. Agents convert these into a strict ordinal ranking over items, breaking any ties in utility in some fixed way.

The welfare of an agent is simply the sum of the utilities of the items allocated to that agent. The utilitarian welfare is the sum of the welfare of the agents, whilst the egalitarian welfare is that of the worst off agent (or agents). The sequential allocation mechanism is parameterized by the policy, the order in which agents pick items. For example, with the policy 123321, agent 1 picks first, then agent 2, then agent 3 before we repeat in reverse. An allocation is an assignment of items to agents. One allocation *Pareto improves* another iff each agent has at least the same utility in the first, and there is at least one agent where the utility is greater. An allocation is *Pareto efficient* iff there is no allocation which Pareto improves it. For every Pareto efficient allocation, there exists a policy such that sincere picking with this policy generates this allocation. We can construct this policy using the greedy algorithm in the proof of Proposition 1 in [8]. The reverse, however, is not true. Sincere picking may not return a Pareto efficient allocation.

Remark 1. *Sincere picking can generate allocations that are not*

Pareto efficient.

Proof: Consider the policy 1221. Suppose the agents' utilities are as follows

	a	b	c	d
1	5	4	2	0
2	8	2	1	0

Both agents have the same total utility over the items. Sincere picking gives items a and d to agent 1 and items b and c to agent 2. This gives an utility of 5 to agent 1 and of 3 to agent 2. If they swap allocations, then the utility of agent 1 increases to 6, and of agent 2 to 8. Hence, sincere picking leads to an allocation that is not Pareto efficient, and does not have the optimal egalitarian or utilitarian social welfare. \diamond

We contrast this observation with Proposition 1 in [8]. This looks just at the rank of items in an agent's preference ordering, ignoring their precise utilities. Given two sets of items S and S' with $|S| = |S'|$, an allocation of items S to an agent *dominates* the allocation of items S' iff for every item in $S - S'$ there is a different item in $S' - S$ that is strictly less preferred. They then define an ordering, *ordinal efficiency* in terms of such domination. This is a strictly weaker ordering than Pareto efficiency which is defined in terms of utilities rather than ordinal rankings.

Proposition 1 in [8] demonstrates that ordinal efficiency corresponds *exactly* to allocations generated by sequential allocation supposing sincere picking. On the other hand, only a subset of the allocations returned by sequential allocation are Pareto efficient, and only a subset again maximize the egalitarian social welfare. However, one of these allocations is certainly Pareto efficient.

Remark 2. *There exists an allocation with the maximum possible egalitarian social welfare that is also Pareto efficient.*

It follows quickly that there always exists a policy for sequential allocation that gives an allocation with the maximum possible egalitarian social welfare supposing sincere picking. Note that this does not rule out other allocations which maximize egalitarian social welfare which are not ordinal efficient, and which cannot be generated by sequential allocation with sincere picking.

Example 1. *Suppose we have three agents (1 to 3), three items (a to c), and Borda utilities. Let agent 1 have a preference order bac , agent 2 have abc , and agent 3 have acb . Then allocating a to 1, b to 2 and c to 3 maximizes the egalitarian social welfare. However, there is no policy for sequential allocation that will return such an allocation supposing agents pick sincerely as no agent gets a first choice item.*

Maximizing the utilitarian social welfare also does not conflict with Pareto efficiency. In this case, we point out the well-known fact that *any* allocation that maximizes utilitarian social welfare is Pareto efficient.

Remark 3. *Any allocation with the maximum possible utilitarian social welfare is also Pareto efficient.*

Again it follows quickly that there exists a policy for sequential allocation that gives an allocation with the maximum possible utilitarian social welfare supposing sincere picking.

POSSIBLE AND NECESSARY WELFARE

Since sequential allocation may not return allocations that are optimal from either an egalitarian or utilitarian perspective, we turn to the (computational) questions of what social welfare is possible or

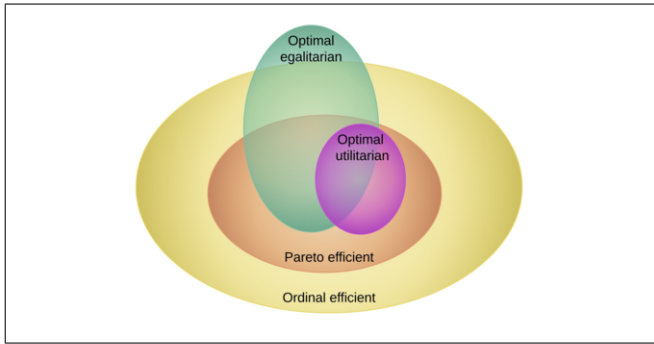


Figure 1: Inclusion relationship between different allocations. Sequential allocation with sincere picking can generate any ordinally efficient allocation.

necessary. Note that throughout this paper, we suppose agents pick sincerely. Whilst strategic behaviour may be beneficial, risk averse agents will tend to pick sincerely, especially when the policy and/or utilities are private information. Nevertheless, it is interesting future work to consider agents acting strategically [14]. We consider four decision problems related to the egalitarian or utilitarian welfare possible or necessary. We also consider different class of sequential mechanisms that depend on the class of the picking sequences allowed.

POSSIBLE EGALITARIAN WELFARE

Input: a set of n items, m agents each with utilities over the items, a class of policies, and an integer t .

Question: Is there a policy in the class that results in an allocation with an egalitarian social welfare of t or greater supposing agents pick items sincerely?

POSSIBLE UTILITARIAN WELFARE

Input: a set of n items, m agents each with utilities over the items, a class of policies, and an integer t .

Question: Is there a policy in the class that results in an allocation with an utilitarian social welfare of t or greater supposing agents pick items sincerely?

NECESSARY EGALITARIAN WELFARE

Input: a set of n items, m agents each with utilities over the items, a class of policies, and an integer t .

Question: Does every policy in the class result in an allocation with an egalitarian social welfare of t or greater supposing agents pick items sincerely?

NECESSARY UTILITARIAN WELFARE

Input: a set of n items, m agents each with utilities over the items, a class of policies, and an integer t .

Question: Does every policy in the class result in an allocation with an utilitarian social welfare of t or greater supposing agents pick items sincerely?

The possible and necessary welfare questions answer a *policy control* problem: can the chair choose a policy to achieve a given social welfare? Similar control problems have been considered previously [1] but with the goal of allocating particular items to agents, rather than, as here, of achieving a particular welfare. Note that we suppose we know the (private) utilities of the agents. As mentioned before, we may relax this assumption by supposing that the utilities are simple functions of the ordinal rank (e.g. Borda, lexicographical or

quasi-indifferent scores). As this is a special case of general utilities, any result that control takes polynomial time in the general case will map onto a polynomial time result in this more restricted setting. When we prove that a particular possible or necessary welfare problem takes polynomial time to solve, we will typically do so by answering a closely related maximization or minimization problem. Such problems are interesting in their own right. We consider four such function problems that compute the maximal or minimal social welfare.

MINIMUM EGALITARIAN WELFARE

Input: a set of n items, m agents each with utilities over the items, and a class of policies.

Output: The minimum egalitarian social welfare possible over all policies supposing agents pick items sincerely.

MINIMUM UTILITARIAN WELFARE

Input: a set of n items, m agents each with utilities over the items, and a class of policies.

Output: The minimum utilitarian social welfare possible over all policies supposing agents pick items sincerely.

MAXIMUM EGALITARIAN WELFARE

Input: a set of n items, m agents each with utilities over the items, and a class of policies.

Output: The maximum egalitarian social welfare possible over all policies supposing agents pick items sincerely.

MAXIMUM UTILITARIAN WELFARE

Input: a set of n items, m agents each with utilities over the items, and a class of policies.

Output: The maximum utilitarian social welfare possible over all policies supposing agents pick items sincerely.

Throughout the paper, we assume that agents have strict ordinal preferences. In some of the proofs, whenever there are ties in the *utilities*, we can obtain strict ordinal preferences by perturbing the utilities by an arbitrarily small margin. The arguments for the reductions are not affected.

ALL POSSIBLE POLICIES

We first consider the case when any policy, balanced or unbalanced is possible. In this case it is easy to maximize the utilitarian social welfare. The chair just need to choose a policy that gives items to the agents which value them most. Recall that we are assuming throughout that agents are behaving sincerely.

Theorem 1. *The MAXIMUM and POSSIBLE UTILITARIAN WELFARE problems take polynomial time to solve.*

Proof: We order the items by the maximum utility assigned by any agent. Ties can be broken in any way. We then construct the policy that allocates items in this order choosing the agent who gives an item the greater utility. No allocation can do better than this. \diamond

The MINIMUM EGALITARIAN WELFARE problem also takes polynomial time to solve. It is always zero as there are policies in which one agents gets no turns to pick items. On the other hand, the POSSIBLE EGALITARIAN WELFARE problem is intractable in general, even in the special case that all the agents have identical utilities for the items. We note that this problem has previously been shown to be NP-complete but not in the strong sense (Theorem 5.1 in [16]).

Theorem 2. *The POSSIBLE EGALITARIAN WELFARE problem for m items and n agents is strongly NP-complete when $m \geq 2n$ even when agents have the same ordinal preferences.*

Proof: Membership in NP is shown by giving the policy. To show NP-hardness, we consider $m = 2n$. For larger m , we add dummy items to which all agents assign the same zero utility. Recall that we have some fixed tie breaking mechanism to order items with the same utility. In this and other proofs that follow, we can replace items with zero utility by items with some uniform small, non-zero utility. This merely makes the algebra a little more complex so, to ease exposition, we present proofs here using items of zero utility. The proof uses a reduction from numerical 3-dimensional matching. Given an integer t and 3 multisets $X = \{x_1, \dots, x_n\}$, $Y = \{y_1, \dots, y_n\}$ and $Z = \{z_1, \dots, z_n\}$ of integers with $\sum_{i=1}^n (x_i + y_i + z_i) = nt$, this problem asks if there are permutations σ and π such that $x_i + y_{\sigma(i)} + z_{\pi(i)} = t$ for all $i \in [n]$. We construct an allocation problem over n agents and $m \geq 2n$ items as follows. Let $u = 1 + \sum_{i=1}^n z_i$. For every $j \in [n]$, there is a “big” item with utility $u + x_i + y_j$ for agent i ($i = 1, \dots, n$) and a “small” item which all agents give utility z_j . Finally, there are $m - 2n$ items with zero utility for all agents. We ask if we can achieve an egalitarian welfare of $u + t$. To achieve this, each agent must get precisely a utility of $u + t$. This is only possible if each agent gets one big item and one small item, and $x_i + y_{\sigma(i)} + z_{\pi(i)} = t$ where $\sigma(i)$ and $\pi(i)$ denote are the indices of the big and the small item obtained by agent i . Therefore, we can achieve the egalitarian welfare of $u + t$ iff there is a solution of the original numerical 3-dimensional matching problem. \diamond

Computing the maximum egalitarian welfare possible is also intractable in general.

Theorem 3. *The MAXIMUM EGALITARIAN WELFARE problem for m items and n agents is NP-hard to compute when $m \geq 2n$ even when the agents have the same ordinal preferences.*

Proof: We use the same reduction as in the last proof. \diamond

In the more restricted setting that utilities are Borda scores, computing the possible egalitarian welfare remains intractable. We thank an anonymous reviewer of an earlier version of this paper for suggesting this result.

Theorem 4. *With Borda utilities, the POSSIBLE EGALITARIAN WELFARE problem for m items and n agents is NP-complete when $m \geq \frac{12n+4}{5}$.*

Proof: We use the reduction in the proof of Theorem 3 in [4]. This reduction proves that deciding if there is an allocation with an egalitarian social welfare greater than or equal to some constant t is NP-complete even when utilities are Borda scores. Note that, unlike the previous two proofs, agents in this reduction do not share the same ordinal preferences. It is easy to show that there is a policy for sequential allocation that finds the precise allocation constructed in this reduction. Note also that $\frac{12n+4}{5} > 2n$ so this reduction uses slightly more items than the previous two proofs, in addition to requiring agents to have different ordinal preferences. \diamond

It follows that the MAXIMUM EGALITARIAN WELFARE problems is NP-hard to compute in this setting. Similarly, with lexicographical utilities it is intractable to compute the egalitarian welfare possible,

Theorem 5. *With lexicographical utilities, the POSSIBLE EGALITARIAN WELFARE problem for m items and n agents is NP-complete when $m \geq 3n$.*

Proof: This follows from Theorem 1 in [4]. Note that we again need more than the $2n$ items used in the first two reductions. We have also again relaxed the assumption that the agents have the same ordinal preferences. \diamond

It follows immediately that the MAXIMUM EGALITARIAN WELFARE problem is NP-hard to compute with lexicographical utilities. Finally, the NECESSARY EGALITARIAN WELFARE problem is trivial: if an agent does not get a turn, their welfare is zero.

BALANCED POLICIES

It might be considered unfair to use any policy, for example one in which one agent gets many more items than another. Whilst looking for allocations that maximize fairness and efficiency, Brams and King [8] observe that “*the symbolic value of giving players equal numbers of items, such as landing slots at an airport, may be important*”. We therefore consider the restricted class of balanced policies. In a *balanced* policy, each agent gets the same number of items. Recall that we suppose the number of items is an integer multiple of the number of agents. Hence, we can give each agent the same number of items. Of course, limiting sequential allocation to balanced policies impacts the social welfare that can be obtained.

To maximize utilitarian welfare, we cannot simply give items to the agents that value them most. This may violate balance. Despite this restriction, we can still find the policy that maximizes the utilitarian welfare in polynomial time.

Theorem 6. *The MAXIMUM and POSSIBLE UTILITARIAN WELFARE problems for balanced policies take polynomial time to solve.*

Proof: We suppose that there are kn items to divide between the n agents. We set up a min cost max flow problem. We connect the source node to nodes representing the agents, each with a capacity of k and no cost. We connect the nodes representing agents to nodes representing the items. Each edge has a capacity of 1, and a cost equal to minus the utility that the agent assigns to the item. Finally we connect the nodes representing the items to the target node, each with an edge of capacity 1 and zero cost. We find a Pareto efficient allocation from any such flow using the top trading cycle algorithm [17]. A policy can be constructed that achieves this Pareto efficient allocation by again exploiting Proposition 1 in [8]. \diamond

By comparison, the NECESSARY UTILITARIAN WELFARE problem is intractable for balanced policies.

Theorem 7. *The NECESSARY UTILITARIAN WELFARE problem for balanced policies is coNP-complete for $m \geq 2n$.*

Proof: We reduce from the NECESSARY ITEM problem. This asks if a given agent necessarily gets an given item irrespective of the policy used. The NECESSARY ITEM problem for balanced policies is coNP-complete even when limited to an agent’s most preferred item and $m = 2n$ [1]. Let one agent have utility of 1 for her most preferred item, zero utility for all others, and the other agents all have utility 1 for every item. Then the NECESSARY ITEM problem is equivalent to asking if an utilitarian welfare of m or more is necessary. \diamond

It follows that the MINIMUM UTILITARIAN WELFARE problem for balanced policies is NP-hard to compute. Restricting to balanced policies also does not change the intractability of computing the egalitarian welfare that is possible.

Theorem 8. *The POSSIBLE EGALITARIAN WELFARE problem for balanced policies is NP-complete for $m \geq 2n$ even when agents have the same ordinal preferences.*

Proof: This follows almost immediately from the reduction used in the proof of Theorem 2. Note that this reduction uses policies in which not all agents get the same number of items. However, such unbalanced policies can be ignored as they result in poor egalitarian welfare. Note also that when a numerical 3-dimensional matching exists, the corresponding successful policy constructed in the reduction is balanced. \diamond

It follows immediately that MAXIMUM EGALITARIAN WELFARE problem is NP-hard to compute for balanced policies. Note that an easy reduction from the EQUI-PARTITION problem demonstrates that the POSSIBLE EGALITARIAN WELFARE problem for balanced policies is NP-complete even for just two agents with identical utilities. On the one hand, this is a more restricted setting than Theorem 8 as we now have only 2 agents and they have identical utilities. On the other hand, this is a weaker result, as it is not strong NP-completeness, and dynamic programming will return a result in polynomial time supposing utilities are specified in unary.

With lexicographical utilities and balanced policies, it remains intractable to compute the egalitarian welfare possible.

Theorem 9. *With lexicographical utilities, the POSSIBLE EGALITARIAN WELFARE problem for balanced policies is NP-complete for $m \geq 3n$.*

Proof: We can adapt the reduction in the proof of Theorem 1 in [4]. Note that we have again relaxed the assumption from Theorem 8 that agents have the same ordinal preferences. \diamond

Finally, computing the egalitarian welfare necessary is intractable for balanced policies.

Theorem 10. *The NECESSARY EGALITARIAN WELFARE problem for balanced policies is coNP-complete for $m \geq 2n$.*

Proof: The same reduction as in the proof of Theorem 7. \diamond

RECURSIVELY BALANCED POLICIES

Balanced policies might still be considered unfair. For example, a policy like 11112222 favours the first agent even though it is balanced, guaranteed to return a Pareto efficient allocation, and is strategy proof. We therefore consider an even more restrictive class: recursively balanced policies. In such a policy, items are allocated in rounds, and each agent appears once in each round. For simplicity, we again suppose that the number of items is an integer multiple of the number of agents and add dummy items of no utility otherwise. When the number of items equals the number of agents, all balanced policies are recursively balanced. For this reason, we focus on problems where the number of items exceeds the number of agents.

Formally, a policy is *recursively balanced* iff it is the empty policy, or it is non-empty and every agent has exactly one turn in the first n picks and the remaining policy is also recursively balanced. Recursively balanced policies include the alternating policy (12121212...), the balanced alternating policy (12211221...), as well as the Thue-Morse sequence (1221211221...). With two agents, recursively balanced policies are concatenations of 12 and 21. Other simple properties of recursively balanced policies follow immediately from their definition. For example, no agent has more than two successive picks in a recursively balanced policy. Limiting sequential allocation to recursively balanced policies may further impact the social welfare that can be obtained.

There are several situations where focusing on recursively balanced policies does not hurt welfare. For example, with Borda utilities, the expected utilitarian social welfare for two agents is not impacted by limiting allocation to recursively balanced policies. The

simple alternating policy which is recursively balanced is optimal in expectation [13]. Similarly for Borda utilities and small n , the expected egalitarian social welfare for two agents is not impacted. The authors [13] computed the policies that maximize expected egalitarian social welfare for up to 12 items and for each n , at least one optimal policy is recursively balanced.

In general, restricting to recursively balanced policies results in it being intractable to decide if a given egalitarian or utilitarian welfare can or must be achieved.

Theorem 11. *The POSSIBLE EGALITARIAN and POSSIBLE UTILITARIAN WELFARE problems for recursively balanced policies are NP-complete for $m \geq 2n$, whilst the NECESSARY EGALITARIAN and NECESSARY UTILITARIAN WELFARE are coNP-complete for $m \geq 3n$.*

Proof: We reduce from the corresponding problem of deciding whether the top k most preferred items of an agent are possible or necessary [1]. TOP- k POSSIBLE SET for recursively balanced policies is NP-complete for $m \geq 3n$. We reduce this to POSSIBLE EGALITARIAN WELFARE as follows. Let one agent have utility of k^2 for their i th most preferred items ($i \leq k$), zero utility for all others, and the other agents all have utility k^3 or greater for any item. Then TOP- k POSSIBLE SET is equivalent to asking if an egalitarian welfare of k^3 or more is possible. We also reduce TOP- k POSSIBLE SET to POSSIBLE UTILITARIAN WELFARE as follows. Let one agent have utility of mk^2 for their i th most preferred items ($i \leq k$), zero utility for all others, and all other agents have utility of k or less for any item. Then TOP- k POSSIBLE SET is equivalent to asking if an utilitarian welfare of mk^3 or more is possible.

TOP- k NECESSARY SET is coNP-complete for recursively balanced policies when $m \geq 2n$. We reduce this to NECESSARY EGALITARIAN WELFARE as follows. Let one agent have total utility of k^2 for their k most preferred items zero utility for all others, and the other agents all have utility k^3 or greater for any item. Then TOP- k NECESSARY SET is equivalent to asking if an egalitarian welfare of k^2 is necessary. We also reduce TOP- k NECESSARY SET to NECESSARY UTILITARIAN WELFARE as follows. Let one agent have utility of mk^2 for their i th most preferred items ($i \leq k$), zero utility for all others, and all other agents have utility of k or less for any item. Then TOP- k NECESSARY SET is equivalent to asking if an utilitarian welfare of mk^3 or more is necessary. \diamond

Even when agents have identical utilities, these problems can remain intractable.

Theorem 12. *When allocating $2n$ items between two agents, the POSSIBLE EGALITARIAN WELFARE problem for recursively balanced policies is NP-complete even when agents have identical utilities given in binary.*

Proof: Membership in NP is clear. For the hardness we use reduction from PARTITION: for positive integers a_1, \dots, a_n with $a_1 + \dots + a_n = 2B$, the problem is to decide if there is a nonempty set $I \subseteq [n]$ with $\sum_{i \in I} a_i = B$. We reduce this to the POSSIBLE EGALITARIAN WELFARE problem for two agents and $2n$ items with utilities $c_1 = 2B, c_{2n} = 0$, and

$$c_{2j} = c_{2j+1} = c_{2j-1} - a_j \quad \text{for } j = 1, 2, \dots, n-1.$$

Let $C = \sum_{i=1}^{2n} c_i$ be the sum of the utilities, and u_i be the utility received by agent i in a given allocation. Note that an egalitarian welfare of $C/2$ is equivalent to $u_1 = u_2$. In round j , the items with utilities c_{2j-1} and c_{2j} are allocated. From $c_{2j-1} - c_{2j} = a_j$ it follows

that the difference $u_1 - u_2$ between the agents' utilities increases by a_j if agent 1 starts and decreases if agent 2 starts. Let $I \subseteq [n]$ be the set of rounds in which agent 1 starts. An egalitarian social welfare of $C/2$ is achieved if and only if

$$0 = u_1 - u_2 = \sum_{j \in I} a_j - \sum_{j \in [n] \setminus I} a_j,$$

That is, if and only if there is a perfect partition. \diamond

BALANCED ALTERNATING POLICIES

The final and most restricted class of policies we consider is that of balanced alternating. This is the subclass of recursively balanced policies in which each round is the reverse of the previous. When allocating students to courses at the Harvard Business School, such a policy is chosen uniformly at random from the space of all possible balanced alternating policies. This gives a form of procedural fairness.

Theorem 13. *The POSSIBLE EGALITARIAN and POSSIBLE UTILITARIAN WELFARE problems for balanced alternating policies are NP-complete for $m \geq 2n$, whilst the NECESSARY EGALITARIAN and NECESSARY UTILITARIAN WELFARE are coNP-complete again for $m \geq 2n$.*

Proof: As in proof of Theorem 11, by reduction from corresponding TOP- k POSSIBLE or NECESSARY SET problem for balanced alternating policies. Given an allocation problem, preference profiles for all the agents, a class of policies, and a designated agent, the TOP- k POSSIBLE problem asks if there is a policy such that the agent gets their top k most preferred items. Given an allocation problem, preference profiles for all the agents, a class of policies, and a designated agent, the TOP- k NECESSARY SET problem asks if the agent necessarily gets their top k most preferred items irrespective of the policy used. TOP- k POSSIBLE for balanced alternating policies is NP-complete for $m \geq 2n$, whilst TOP- k NECESSARY SET is coNP-complete [1]. \diamond

It follows that it is NP-hard to compute the probability that the Harvard Business School course allocation mechanism returns an allocation with egalitarian or utilitarian welfare greater than or equal to some given value t .

TWO AGENTS

We now consider some special and more tractable cases. With two agents, we can find a balanced policy that maximizes the egalitarian or utilitarian welfare in polynomial time.

Theorem 14. *The MAXIMUM EGALITARIAN and MAXIMUM UTILITARIAN WELFARE problems with balanced policies can be solved in $O(w^2n^3)$ and $O(wn^2)$ time respectively when allocating $2n$ items between two agents with utilities that are (possibly different) integers taken from $[0, w]$.*

Proof: We put the items into some (arbitrary) order and consider how each item is allocated in turn. We construct a $2n$ step dynamic program in which the i th step corresponds to the decision of where to allocate the i th item in this order. The states of this dynamic program are triples containing the number of items allocated to the first agent, the sum of the utilities of the items so far allocated to the first agent, and the sum of the utilities of the items so far allocated to the second agent. We can compute the number of items allocated to the

second agent from this. As both sums are bounded in size by $2wn$, this dynamic program has $O(w^2n^3)$ states. For the maximum utilitarian welfare, the states of the dynamic program can be simpler and just need to be pairs containing the number of items allocated to the first agent, and the sum of the utilities of the items so far allocated to both agents. \diamond

This result generalizes to a bounded number of agents. On the other hand, when utilities are specified in binary, an easy reduction from the PARTITION problem demonstrates that the POSSIBLE EGALITARIAN WELFARE problem is NP-complete even when the two agents have identical utilities. This is almost identical to Proposition 2 in [5] which shows that deciding if there is a policy that ensures a given expected egalitarian welfare is NP-complete when the utilities of the two agents are identical.

With recursively balanced policies, we consider the case where agents have the same ordinal ranking over items.

Theorem 15. *The MAXIMUM and POSSIBLE EGALITARIAN WELFARE problems for recursively balanced policies can be solved in $O(w^2n^2)$, whilst the MAXIMUM and POSSIBLE UTILITARIAN WELFARE problems can be solved in just $O(wn)$ time when allocating $2n$ items between two agents when agents have the same ordering over items but possibly different utilities, and utilities are integers drawn from $[0, w]$.*

Proof: We construct a n step dynamic program in which each step corresponds to allocating one item to each of the agents. The states of this program are pairs containing the sums of the utilities of items so far allocated to the two agents. As both sums are bounded by wn , there are $O(w^2n^2)$ states. By scanning the final step of the dynamic program that computes all possible partitions, we can compute the optimal egalitarian social welfare. To compute the optimal utilitarian social welfare, we can use a simpler dynamic program where the states are just the sum of the utilities allocated to the agents. \diamond

This result again generalizes to a bounded number of agents easily.

HOUSE ALLOCATION

Another more tractable case is house allocation, when we have only as many items as agents and each agent gets one item. Results for this setting imply results for recursively balanced and balanced alternating policies for $m = n$. For example, sequential allocation is used in many universities and residential colleges to assign rooms to students. In this case, we can solve the MAXIMUM and POSSIBLE EGALITARIAN WELFARE problems over all possible policies in polynomial time. We construct a graph between agents and items with edges for all items that have a utility greater than or equal to the desired egalitarian social welfare. The POSSIBLE EGALITARIAN WELFARE problem is solvable if we can find a perfect matching in this graph. To construct a satisfying policy, we find a Pareto efficient allocation from this matching using the top trading cycle algorithm [17]. A policy can be constructed that achieves this Pareto efficient allocation using Proposition 1 in [8].

We also show that NECESSARY EGALITARIAN WELFARE is polynomial-time solvable for house allocation.

Theorem 16. *NECESSARY EGALITARIAN WELFARE is polynomial-time solvable for house allocation.*

Proof. We first show that it can be checked in polynomial time whether $n - 1$ agents get allocated a target $n - 1$ set of items in the first $n - 1$ turns. The problem reduces to checking whether there

	all policies	balanced	recursively balanced	balanced alternating
POSSIBLE EGALITARIAN WELFARE	NPC P	NPC P (for $m = n$)	NPC P (for $m = n$)	NPC P (for $m = n$)
POSSIBLE UTILITARIAN WELFARE	P P	P P (for $m = n$)	NPC P (for $m = n$)	NPC P (for $m = n$)
NECESSARY EGALITARIAN WELFARE	P P	coNPC P (for $m = n$)	coNPC P (for $m = n$)	coNPC P (for $m = n$)
NECESSARY UTILITARIAN WELFARE	?	coNPC	coNPC	coNPC

Table 1. Summary of results when allocating m items between n agents: NPC=NP-complete, coNPC=coNP-complete, P=polynomial-time. Note that over half the entries in this table identify problems which can be solved in polynomial time.

exists a matching in which each of the $n - 1$ agents gets a more preferred item than the one not in the target set. If the allocation of the n agents is not ordinally efficient, it can be made ordinally efficient via trading cycles none of which will involve the n -th agent's item. This means that we can check whether there exists a sequence such that the n -th agent can get a particular item in the last turn. In order to get minimum egalitarian welfare, some agent i has to get a low preferred item o . If there exists a policy in which i gets o in some turn that is not the last turn, then i gets o or an even worse item if i 's turn is moved to the end of same policy. Thus we can check in polynomial time for each i and o whether i gets o in the last turn and identify the i 's and o 's for which this is possible and $u_i(o)$ is minimum. \square

The same statement also applies to recursively balanced and balanced alternative policies provided that $m = n$.

OTHER RELATED WORK

As mentioned earlier, Bouveret and Lang [5] consider the case in which the utilities of items are simply functions of the ordinal rankings. They prove that any recursively balanced policy tends to an allocation giving the optimal expected egalitarian or utilitarian welfare as the number of items grows, supposing sincere picking, utilities that are Borda scores and all ordinal rankings being equiprobable. In addition, they compute the optimal policies for maximizing the expected egalitarian or utilitarian welfare under the same assumptions for up to 12 items. The optimal policies for two agents and an even number of items are recursively balanced. Kalinowski *et al.* [13] prove that the alternating policy maximizes the expected utilitarian welfare under these same assumptions. Note that such results are about maximizing the *expected* welfare supposing limited knowledge about the utilities, whilst the results here about maximizing the *exact* welfare supposing the chair knows the *actual* utilities.

There has been some study of strategic behaviour of agents (as opposed to the chair) in the sequential allocation mechanism. It can, for example, be viewed as a repeated game. When all agents have complete information, we can compute the subgame perfect Nash equilibrium. This is unique and takes polynomial time to compute for two agents [15, 14], but for an arbitrary number of agents, there can be an exponential number of equilibria and computing even one is PSPACE-hard [14]. More recently, Bouveret and Lang [6] consider how an agent or coalition of agents can strategically mis-report their preferences in a sequential allocation mechanism supposing the other agents act sincerely. They show that the loss of social welfare caused by such manipulation is not great. For example, with Borda scoring, two agents, and the alternating policy, there was at most a 33% loss in the utilitarian welfare.

More recently, a family of rules for dividing indivisible goods among agents has been proposed that take as input the agents' ordinal rankings over the items, a scoring vector, and a social welfare aggregation function [4, 2]. They return the allocation that maximizes the social welfare according to this scoring rule and aggregation function. Whilst such rules have a number of desirable properties like monotonicity, it is computationally challenging to compute the actual allocation (unless we have a bounded number of agents in which case we can typically use dynamic programming). This contrasts with sequential allocation where computing the allocation take just linear time. Baumeister *et al.* [2] also compute the multiplicative/additive "price of elicitation-freeness", the worst-case ratio/difference in social welfare between such allocations and the allocation returned by sequential allocation. Whilst their results are limited to simple alternating policies, the prices are typically not great. For example, the optimal utilitarian welfare with Borda scores is at most twice that returned by sequential allocation using simple alternation.

CONCLUSIONS

We have considered the implications on social welfare of choosing different policies when using a sequential mechanism to allocate indivisible goods. In particular, we consider the (computational) questions of what welfare is possible or necessary. The former is related to the control problem in which a (benevolent) chair chooses a policy for the sequential allocation mechanism to improve the social welfare. These questions are also related to the expected welfare when we choose a policy uniformly at random. Our results are summarized in Table 1. We note again that over half the entries in this table are polynomial time algorithms. Many of these questions are computationally tractable, and our results provide efficient algorithms to answer them. There are many interesting open questions. For example, how difficult is it to find a recursively balanced policy that returns a Pareto efficient allocation supposing sincere agents? As a second example, how can the chair adaptively compute a policy as agents pick items to maximize social welfare? And what is the "price of adaptability", the ratio between the welfare achieved when the chair is adaptive and when the chair has to declare and fix the policy in advance?

REFERENCES

- [1] H. Aziz, T. Walsh, and L. Xia, 'Possible and necessary allocations via sequential mechanisms', in *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, (IJCAI 2015)*, eds., Q. Yang and M. Wooldridge, (2015).

- [2] D. Baumeister, S. Bouveret, J. Lang, N. Nguyen, T. Nguyen, J. Rothe, and A. Saffidine, 'Axiomatic and computational aspects of scoring allocation rules for indivisible goods', in *5th International Workshop on Computational Social Choice (COMSOC)*, pp. 1–22, Pittsburgh, USA, (June 2014).
- [3] D. Baumeister, S. Bouveret, J. Lang, T. Nguyen, N. Nguyen, and J. Rothe, 'Scoring rules for the allocation of indivisible goods', in *Proceedings of 21st European Conference on Artificial Intelligence (ECAI 2014)*, eds., T. Schaub, G. Friedrich, and B. O'Sullivan, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pp. 75–80. IOS Press, (2014).
- [4] D. Baumeister, S. Bouveret, J. Lang, T. Nguyen, J. Rothe, and A. Saffidine, 'Positional scoring rules for the allocation of indivisible goods', in *11th European Workshop on Multi-Agent Systems (EUMAS)*, pp. 1–14, Toulouse, France, (December 2013).
- [5] S. Bouveret and J. Lang, 'A general elicitation-free protocol for allocating indivisible goods', in *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, ed., Toby Walsh, pp. 73–78. IJCAI/AAAI, (2011).
- [6] S. Bouveret and J. Lang, 'Manipulating picking sequences', in *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*. IOS Press, (2014).
- [7] S. Brams, P. Edelman, and P. Fishburn, 'Fair division of indivisible items', *Theory and Decision*, **55**(2), 147–180, (2003).
- [8] S.J. Brams and D.L. King, 'Efficient fair division: Help the worst off or avoid envy?', *Rationality and Society*, **17**(4), 387–421, (2005).
- [9] S.J. Brams and A.D. Taylor, *Fair Division: From cake-cutting to dispute resolution*, Cambridge University Press, Cambridge, 1996.
- [10] E. Budish and E. Cantillon, 'The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard', *American Economic Review*, **102**(5), 2237–71, (2012).
- [11] A. Darmann and J. Schauer, 'Maximizing nash product social welfare in allocating indivisible goods', *European Journal of Operational Research*, **247**(2), 548 – 559, (2015).
- [12] E. Fujita, J. Lesca, A. Sonoda, T. Todo, and M. Yokoo, 'A complexity approach for core-selecting exchange with multiple indivisible goods under lexicographic preferences', in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, eds., B. Bonet and S. Koenig, pp. 907–913. AAAI Press, (2015).
- [13] T. Kalinowski, N. Narodytska, and T. Walsh, 'A social welfare optimal sequential allocation procedure', in *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-2013)*. International Joint Conference on Artificial Intelligence, (2013).
- [14] 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).
- [15] D.A. Kohler and R. Chandrasekaran, 'A Class of Sequential Games', *Operations Research*, **19**, 270–277, (1971).
- [16] M. Roos and J. Rothe, 'Complexity of social welfare optimization in multiagent resource allocation', in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, AAMAS10*, pp. 641–648, (2010).
- [17] L. Shapley and H. Scarf, 'On cores and indivisibility', *Journal of Mathematical Economics*, **1**(1), 23 – 37, (1974).