

A Social Welfare Optimal Sequential Allocation Procedure*

Thomas Kalinowski

Universität Rostock
Rostock, Germany

thomas.kalinowski@uni-rostock.de

Nina Narodytska and Toby Walsh

NICTA and UNSW
Sydney, Australia

{nina.narodytska,toby.walsh}@nicta.com.au

Abstract

We consider a simple sequential allocation procedure for sharing indivisible items between agents in which agents take turns to pick items. Supposing additive utilities and independence between the agents, we show that the expected utility of each agent is computable in polynomial time. Using this result, we prove that the expected utilitarian social welfare is maximized when agents take alternate turns. We also argue that this mechanism remains optimal when agents behave strategically.

1 Introduction

There exist a variety of mechanisms to share indivisible goods between agents without side payments [Brams and Fishburn, 2000; Herreiner and Puppe, 2002; Brams *et al.*, 2003; Brams and Kaplan, 2004; Brams *et al.*, 2012]. One of the simplest is simply to let the agents take turns to pick an item. This mechanism is parameterized by a policy, the order in which agents take turns. In the alternating policy, agents take turns in a fixed order, whilst in the balanced alternating policy, the order of agents reverses each round. Bouveret and Lang (2011) study a simple model of this mechanism in which agents have strict preferences over items, and utilities are additive. They conjecture that computing the expected social welfare for a given policy is NP-hard supposing all preference orderings are equally likely. Based on simulation for up to 12 items, they conjecture that the alternating policy maximizes the expected utilitarian social welfare for Borda utilities, and prove it does so asymptotically. We close both conjectures. Surprisingly, we prove that the expected utility of each agent can be computed in polynomial time for any policy and utility function. Using this result, we prove that the alternating policy maximizes the expected utilitarian social welfare for any number of items and any linear utility function including Borda. Our results provides some justification for a mechanism in use in school playgrounds around the world.

*Supported by the Department of Broadband, Communications and the Digital Economy, the ARC and the Asian Office of Aerospace Research and Development (AOARD-124056).

2 Notation

Each of the n agents has a total preference order over the p items. A *profile* is an n -tuple of such orders. Agents have the same utility function. An item ranked in k th position has a utility $g(k)$. For Borda utilities, $g(k) = p - k + 1$. The utility of a set of items is the sum of their individual utilities. Preference orders are independent and drawn uniformly at random from all $p!$ possibilities (full independence). Agents take turns to pick items according to a *policy*, a sequence $\pi = \pi_1 \dots \pi_p \in \{1, 2, \dots, n\}^p$. At the k -th step, agent π_k chooses one item from the remaining set. Without loss of generality, we suppose $\pi_1 = 1$. For profile R , $u_R(i, \pi)$ denotes the utility gained by agent i supposing every agent always chooses the item remaining that they rank highest. We write $\bar{u}_i(\pi)$ for the expectation of $u_R(i, \pi)$ over all possible profiles. We take an utilitarian standpoint, measuring social welfare by the sum of the utilities: $\text{sw}_R(\pi) = \sum_{i=1}^n u_R(i, \pi)$. By linearity of expectation, $\overline{\text{sw}}(\pi) = \sum_{i=1}^n \bar{u}_i(\pi)$. To help compute expectations, we need a sequence γ_k given by $\gamma_1 = \gamma_2 = 1$,

$$\gamma_k = \prod_{j=1}^{\lfloor (k-1)/2 \rfloor} \frac{2j+1}{2j} \text{ for } k \geq 3 \text{ and } \bar{\gamma}_k = \gamma_k/k \text{ for all } k.$$

Asymptotically $\gamma_k = \sqrt{\frac{2k}{\pi}} + O\left(\frac{1}{\sqrt{k}}\right)$. To simplify notation, we suppose empty sums are zero and empty products are one.

3 Computing the Expected Social Welfare

Bouveret and Lang 2011 conjectured that it is NP-hard to compute the expected social welfare of a given policy as such a calculation depends on the super-exponential number of possible profiles. However, as we show now, the expected utility of each agent can be computed in just $O(np^2)$ time for an arbitrary utility function, and $O(np)$ time for Borda utilities. We begin with this last case, and then extend the results to the general case. Let \mathcal{P}_p^n denote the set of all policies of length p for n agents. For $p \geq 2$, we define an operator $\mathcal{P}_p^n \rightarrow \mathcal{P}_{p-1}^n$ mapping $\pi \mapsto \tilde{\pi}$, by deleting the the first entry. More precisely, $\tilde{\pi}_i = \pi_{i+1}$ for $i \in \{1, \dots, p-1\}$. For example, $\pi = 1211$ and $\tilde{\pi} = 211$.

Lemma 1. *For Borda scoring, $n \geq 2$ agents, $p \geq 2$ items and $\pi \in \mathcal{P}_p^n$ with $\pi_1 = 1$, we have:*

$$\bar{u}_1(\pi) = p + \bar{u}_1(\tilde{\pi}), \quad \bar{u}_i(\pi) = \frac{p+1}{p} \bar{u}_i(\tilde{\pi}), i \neq 1$$

and these values can be computed in $O(np)$ time.

Proof. Agent 1 picks her first item, giving her a utility of p . After that, from her perspective, it's the standard game on $p - 1$ items with policy $\tilde{\pi}$, so she expects to get an utility of $\bar{u}_1(\tilde{\pi})$. This proves the first equation. For the other agents, it is more involved. Let $i \in \{2, \dots, n\}$ be a fixed agent. For $q \in \{1, \dots, p\}$, let $a_i(q, \pi)$ denote the probability that under policy π agent i gets the item with utility q . Note that this probability does not depend on the utility function but only on the ranking: it is the probability that agent i gets the item of rank $p - q + 1$ in her preference order. By the definition of expectation,

$$\bar{u}_i(\pi) = \sum_{q=1}^p a_i(q, \pi)q. \quad (1)$$

There are three possible outcomes of the first move of agent 1 with respect to the item that has utility q for agent i . With probability $(q - 1)/p$, agent 1 has picked an item with utility less than q (for agent i), with probability $(p - q)/p$, agent 1 has picked an item with utility more than q , and with probability $1/p$ it was the item of utility equal to q . In the first case there are only $q - 2$ items of utility less than q left, hence the probability for agent i to get the item of utility q is $a_i(q - 1, \tilde{\pi})$. In the second case there are still $q - 1$ items of value less than q , hence the probability to get the item of utility q is $a_i(q, \tilde{\pi})$. In the third case, the probability to get the item of utility is zero, and together we obtain

$$a(q, \pi) = \frac{q - 1}{p} a_i(q - 1, \tilde{\pi}) + \frac{p - q}{p} a_i(q, \tilde{\pi}). \quad (2)$$

Substituting this into (1) yields

$$\begin{aligned} \bar{u}_i(\pi) &= \sum_{q=1}^p \left[\frac{q - 1}{p} a_i(q - 1, \tilde{\pi}) + \frac{p - q}{p} a_i(q, \tilde{\pi}) \right] q \\ &= \sum_{q=1}^p \frac{(q - 1)q}{p} a_i(q - 1, \tilde{\pi}) + \sum_{q=1}^p \frac{(p - q)q}{p} a_i(q, \tilde{\pi}) \end{aligned}$$

In the first sum we substitute $q' = q - 1$ and this yields

$$\bar{u}_i(\pi) = \sum_{q'=0}^{p-1} \frac{q'}{p} \cdot a_i(q', \tilde{\pi}) \cdot (q' + 1) + \sum_{q=1}^p \frac{p - q}{p} \cdot a_i(q, \tilde{\pi}) \cdot q$$

The first term in the first sum and the last term in the second sum are equal to zero, so they can be omitted and we obtain

$$\begin{aligned} \bar{u}_2(\pi) &= \sum_{q'=1}^{p-1} \frac{q'}{p} \cdot a_i(q', \tilde{\pi}) \cdot (q' + 1) + \sum_{q=1}^{p-1} \frac{p - q}{p} \cdot a_i(q, \tilde{\pi}) \cdot q \\ &= \sum_{q=1}^{p-1} a_i(q, \tilde{\pi}) \left[\frac{q}{p} \cdot (q + 1) + \frac{p - q}{p} \cdot q \right] \\ &= \frac{p + 1}{p} \sum_{q=1}^{p-1} a_i(q, \tilde{\pi}) \cdot q = \frac{p + 1}{p} \bar{u}_2(\tilde{\pi}). \end{aligned}$$

The time complexity follows immediately from the recursions. \square

		$\pi^1 = 121212$			$\pi^2 = 111222$				
		π	$\bar{u}_1(\pi)$	$\bar{u}_2(\pi)$	$\overline{sw}(\pi)$	π	$\bar{u}_1(\pi)$	$\bar{u}_2(\pi)$	$\overline{sw}(\pi)$
1	2	0	1	1	2	0	1	1	
2	12	2	1.5	3.5	22	0	3	3	
3	212	2.67	4.5	7.17	222	0	6	6	
4	1212	6.67	5.63	12.3	1222	4	7.5	11.5	
5	21212	8	10.63	18.63	11222	9	9	18	
6	121212	14	12.4	26.4	111222	15	10.5	25.5	

Table 1: Expected utilities and expected utilitarian social welfare computation for $\pi^1 = 121212$ and $\pi^2 = 111222$

Example 1. Consider two agents with Borda utilities and the policies $\pi^1 = 121212$ and $\pi^2 = 111222$. We compute expected utilities and expected social welfare for each of them using Lemma 1. Table 1 shows results up to two decimal places. Note that expected values computed in all examples in the paper coincide with the results obtained by the brute-force search algorithm from [Bouveret and Lang, 2011].

Due to the linearity of Borda scoring the probabilities $a_i(q, \pi)$ in the proof of Lemma 1 cancel, and this will allow us to solve recursions explicitly and to prove our main result about the optimal policy for Borda scoring in Section 5.

In the general case, we can still compute the expected utilities $\bar{u}(i, \pi)$, and thus $\overline{sw}(\pi)$, but we need the probabilities $a_i(q, \pi)$ from the proof of Lemma 1: $a_i(q, \pi)$ is the probability that under policy π , agent i gets the item ranked at position $p - q + 1$ in her preference order. Computing these probabilities using (2) adds a factor of p to the runtime.

Lemma 2. For $n \geq 2$ agents, $p \geq 2$ items, a policy $\pi \in \mathcal{P}_p^n$ and an arbitrary scoring function g , the expected utility for agent i is

$$\bar{u}_i(\pi) = \sum_{q=1}^p a_i(q, \pi)g(q)$$

and can be computed in $O(np^2)$ time.

Lemma 2 allows us to resolve an open question from [Bouveret and Lang, 2011].

Corollary 1. For n agents and an arbitrary scoring utility function g , the expected utility of each agent, as well as the expected utilitarian social welfare can be computed in polynomial time.

For some special policies, the recursions in Lemma 1 can be solved explicitly. A particularly interesting policy is the strictly alternating one (denoted ALTPOLICY) $\pi = 123 \dots n123 \dots n123 \dots n \dots$.

Proposition 1. Let π be the strictly alternating policy of length p starting with 1. The expected utilities and utilitarian social welfare for two agents and Borda scoring are

$$\begin{aligned} \overline{sw}(\pi) &= \frac{1}{3} [(2p - 1)(p + 1) + \gamma_{p+1}] \\ (\bar{u}_1(\pi), \bar{u}_2(\pi)) &= \begin{cases} \left(\frac{p(p+1)}{3}, \frac{p^2-1}{3} + \frac{1}{3}\gamma_{p+1} \right) & \text{if } p \text{ is even,} \\ \left(\frac{p(p+1)}{3} + \frac{1}{3}\gamma_{p+1}, \frac{p^2-1}{3} \right) & \text{if } p \text{ is odd.} \end{cases} \end{aligned}$$

For n agents the expectations are, for all $i \in \{1, \dots, n\}$,

$$\bar{u}_i(\pi) = \frac{p^2}{n + 1} + O(p), \quad \overline{sw}(\pi) = \frac{np^2}{n + 1} + O(p).$$

Proof. A proof for the two agents case can be done by straightforward induction using the recursions from Lemma 1. For n agents we outline the proof. The full proof is given in Appendix A in the online version [Kalinowski *et al.*, 2013a]. First, we solve the recursions for the expected utility of agent i when the number of items is $p \equiv i - 1 \pmod{n}$. Using these values, we approximate the expected utility for the remaining combinations of p and i . \square

Example 2. Consider the policy π^1 from Example 1. Using Proposition 1, we get: $\bar{u}_1(\pi) = \frac{6(6+1)}{3} = 14$, $\bar{u}_2(\pi) = \frac{6^2-1}{3} + \frac{1}{3}\gamma_{6+1} = 11.67 + 2.19/3 = 12.4$ and $\overline{\text{sw}}(\pi) = 26.4$. These values coincide with results in Table 1.

For a fixed number n of agents, we write the number of items as $p = kn + r$ with $0 \leq r < n$. We call a policy $\pi = \pi_1 \dots \pi_p$ *balanced* if $\{\pi_{in+1}, \dots, \pi_{in+n}\} = \{1, 2, \dots, n\}$ for all $i \in \{0, \dots, k-1\}$ and $|\{\pi_{kn+1}, \dots, \pi_{kn+r}\}| = r$. For any balanced policy, the expected utility of any agent lies between that of agents 1 and n in the alternating policy. Thus every agent has expected utility $p^2/(n+1) + O(p)$.

4 Comparison with Other Mechanisms

We compare with two other allocation mechanisms that provide insight into the efficiency of the alternating policy. The best preference mechanism (BESTPREF) allocates each item to the agent who prefers it most breaking ties at random. The random mechanism (RANDOM) allocates items by flipping a coin for each individual item.

Proposition 2. For two agents, and Borda utilities, the expected utilitarian social welfare of BESTPREF is $\frac{(p+1)(4p-1)}{6}$. For n agents, it is $\frac{np^2}{n+1} + \frac{p}{2} + O(1)$.

Proof. Due to space limitations we only present a proof for two agents. The proof for n agents is again given in the online Appendix B. Let S_p be the set of all permutations of $\{1, 2, \dots, p\}$. Then the expected utilitarian social welfare is

$$\frac{1}{p!} \sum_{a \in S_p} \sum_{i=1}^p \max\{i, a_i\} = \frac{1}{p!} \sum_{i=1}^p \sum_{a \in S_p} \max\{i, a_i\}.$$

We split the inner sum into two sums: one for all permutations $a = (a_1, \dots, a_p)$ with $a_i \leq i$, and one for the remaining permutations. Hence, we get

$$\frac{1}{p!} \sum_{i=1}^p \left[\sum_{a \in S_p: a_i \leq i} i + \sum_{a \in S_p: a_i > i} a_i \right].$$

We compute the value of each inner sum separately. In the first sum each term equals i , so we have to determine the number of terms. For a_i there are i possible values $1, 2, \dots, i$, and for a fixed value of a_i there are $(p-1)!$ permutations of the remaining values. So the first sum has $(p-1)!i$ terms of value i , hence it equals $(p-1)!i^2$. The second sum contains for each $j \in \{i+1, \dots, p\}$ exactly $(p-1)!$ terms of value $a_i = j$, hence it equals $(p-1)! \sum_{j=i+1}^p j =$

$\overline{\text{sw}}(\pi)$		
ALTPOLICY	BESTPREF	RANDOM
$\frac{np^2}{n+1} + O(p)$	$\frac{np^2}{n+1} + \frac{p}{2} + O(1)$	$\frac{p^2+p}{2}$

Table 2: Expected utilitarian social welfare for different mechanisms.

$(p-1)!(p(p+1)/2 - i(i+1)/2)$. So the expected utilitarian social welfare is

$$\begin{aligned} & \frac{1}{p!} \sum_{i=1}^p (p-1)! \left[i^2 + \frac{1}{2}p(p+1) - \frac{1}{2}i(i+1) \right] \\ &= \frac{1}{2p} \sum_{i=1}^p [p(p+1) + i^2 - i] = \frac{(p+1)(4p-1)}{6}. \quad \square \end{aligned}$$

Proposition 3. For n agents, and Borda utilities, the expected utilitarian social welfare of RANDOM is $\frac{p(p+1)}{2}$.

Proof. As the probability of each agent obtaining the i th item is $1/n$, the expected utilitarian social welfare is $n \sum_{i=1}^p \frac{i}{n} = \frac{p(p+1)}{2}$. \square

Table 2 summarizes the expected utilitarian social welfares for these mechanisms. Clearly, BESTPREF is an upper bound on the expected utilitarian social welfare for any allocation mechanism. As in [Bouveret and Lang, 2011], we define asymptotic optimality of a sequence of policies $(\pi^{(p)})_{p=1,2,\dots}$ where $\pi^{(p)}$ is a policy for p items by

$$\lim_{p \rightarrow \infty} \frac{\overline{\text{sw}}(\pi^{(p)})}{\max_{\pi \in \mathcal{P}_p} \overline{\text{sw}}(\pi)} = 1.$$

As can be seen from the table, ALTPOLICY is an asymptotically optimal policy. By the observation in the end of the previous section the same is true for any balanced policy, and this implies Proposition 5 in [Bouveret and Lang, 2011]. However, the proof in [Bouveret and Lang, 2011] is incorrect as it implies that the expected utility is $p^2/n + O(1)$ for every agent which contradicts our upper bound for BESTPREF. See Appendix C for a detailed discussion of the gaps in the proof. Of course, for any given p and preference orderings, ALTPOLICY may not give the maximal utilitarian social welfare possible.

Example 3. Consider two agents and six items with the following preferences: $1 > 2 > 3 > 4 > 5 > 6$ and $1 > 6 > 2 > 3 > 4 > 5$. The ALTPOLICY policy gives items $\{1, 2, 4\}$ and $\{6, 3, 5\}$ to agents 1 and 2 respectively. Hence, the total welfare is $(6+5+3) + (5+3+1) = 23$. Consider a policy $\pi = 121111$ which gives the following items to agents: $\{1, 2, 3, 4, 5\}$ and $\{6\}$. The total welfare is now $(6+5+4+3+2) + (6) = 25$.

The RANDOM mechanism gives the worst expected utilitarian social welfare among the three mechanisms. Moreover, as n increases the expected utilitarian social welfare produced by RANDOM declines compared with the other two mechanisms: $\lim_{p \rightarrow \infty} \frac{\overline{\text{sw}}(\text{RANDOM}^{(p)})}{\overline{\text{sw}}(\text{BESTPREF}^{(p)})} = \frac{n+1}{2n}$.

With two agents, the expected loss using ALTPOLICY compared to BESTPREF (which requires full revelation of the

preference orders) is less than $p/6$. In particular, with high probability ALTPOLICY yields an utilitarian social welfare very close to the upper bound.

Proposition 4. *For two agents and any $\varepsilon > 0$, with probability at least $1 - \varepsilon$, ALTPOLICY is a $(1 - \frac{1}{3p\varepsilon})$ -approximation of the optimal expected utilitarian social welfare.*

Proof. Let the random variables z_1 and z_2 denote the utilitarian social welfare for ALTPOLICY and BESTPREF. Then $z_2 - z_1 \geq 0$ and for the expectations we have $\mathbf{E}(z_1) = \frac{2p^2+p-1+\gamma_p}{3} > \frac{2p^2+p}{3}$ and $\mathbf{E}(z_2) = \frac{4p^2+3p-1}{6}$.

So $\mathbf{E}(z_2 - z_1) < p/6$, and by Markov's inequality

$$\mathbf{P}\left(z_2 - z_1 \geq \frac{p}{6\varepsilon}\right) \leq \frac{p/6}{p/(6\varepsilon)} = \varepsilon.$$

Writing it multiplicatively, with probability at least $1 - \varepsilon$,

$$\frac{z_1}{z_2} > \frac{p^2/2 - p/(6\varepsilon)}{p^2/2} = 1 - \frac{1}{3p\varepsilon}. \quad \square$$

A similar result holds for more than two agents.

Proposition 5. *For n agents, there exists a constant C such that for every $\varepsilon > 0$ with probability at least $1 - \varepsilon$ ALTPOLICY is a $(1 - \frac{C}{p\varepsilon})$ -approximation of the optimal expected utilitarian social welfare.*

5 Optimality of the Alternating Policy

We now consider the problem of finding the policy that maximizes the expected utilitarian social welfare for Borda utilities. Bouveret and Lang 2011 stated that this is an open question, and conjectured that this problem is NP-hard. We close this problem, by proving that ALTPOLICY is in fact the optimal policy for any given p with two agents.

Theorem 1. *The expected utilitarian social welfare is maximized by the alternating policy for two agents supposing Borda utilities and the full independence assumption.*

Note that by linearity of expectation this implies optimality of the alternating policy for every linear scoring function $g(k) = \alpha k + \beta$ with $\alpha, \beta \in \mathbb{R}$, $\alpha \leq 0$. In particular, the result also holds for quasi-indifferent scoring where $g(k) = N + (p - k + 1)$ for large N .

In the following let π_p^* always be the alternating policy of length p . We also recall that due to symmetry we can only consider policies that starts with 1, e.g. policy 212 is equivalent to 121. To prove Theorem 1 we need to prove that for any policy π of length p the expected utilitarian social welfare is smaller or equal to the expected utilitarian social welfare of π_p^* . That is, $\text{dsw}_\pi = \overline{\text{sw}}(\pi) - \overline{\text{sw}}(\pi_p^*) \leq 0$. We proceed in two steps. First, we describe dsw_π recursively, by representing the policy π in terms of its deviations from π_p^* . Second, given the recursive description of dsw_π , we prove by induction that this difference is never positive (Proposition 6). The proof is not trivial as the natural inductive approach to derive $\text{dsw}_\pi \leq 0$ from $\text{dsw}_{\tilde{\pi}} \leq 0$ does not go through. Hence, we will prove a stronger result in Theorem 2 that implies Proposition 6 and Theorem 1.

Recursive definition. To obtain a recursive definition of dsw_π , we observe that any policy π can be written in terms of its deviations from ALTPOLICY policy π^* . We explain this idea using the following example. Consider a policy $\pi = 1121$. There are two ways to extend π with a prefix to obtain policies of length 5: $\pi' = 11121$ and $\pi'' = 21121$ which is equivalent to $\pi'' = 12212$. We say that $\pi'' = 12212$ follows ALTPOLICY in extending π as its prefix is (12) which coincides with the alternation step. We say that $\pi' = 11121$ deviates from ALTPOLICY in extending π as its prefix is (11) which does not correspond to the alternation step.

Next we define a notion of *the policy tree*, which is a balanced binary tree, that represents all possible policies in terms of deviations from π^* . The main purpose of this notion is to explain intuitions behind our derivations and proofs. We start with the policy (1), which is the root of the tree. We expand a policy to the left by a prefix of length one. We can follow the strictly alternation policy by expanding (1) with prefix 2. This gives policy (21) which is equivalent to (12) due to symmetry. Alternatively, we can deviate from ALTPOLICY by expanding (1) with prefix 1. This gives policy (11). This way we obtain all policies of length 2. We can continue expanding the tree from (12) and (11) following the same procedure and keeping in mind that we break symmetries by remembering only policies that start with 1. The following example show all policies of length at most 5. By convention, given a policy π in a node of the tree we say that we follow ALTPOLICY on the left branch and deviate from ALTPOLICY on the right branch.

Example 4. *Figure 1 shows a tree which represents all policies of length at most 5. A number below each policy shows the value of the expected utilitarian social welfare for this policy. As can be seen from the tree, ALTPOLICY is the optimum policy for all p . Consider, for example, $\pi = 12212$. We can obtain this by deviations from π_5^* (shown as the dashed path): (1) \rightarrow_{L_1} (12) \rightarrow_{L_2} (121) \rightarrow_{R_3} (1121) \rightarrow_{L_4} (12212).*

Next we give a formal recursive definition of dsw_π . We recall that from Lemma 1 the recursions for ALTPOLICY

$$(\overline{u}_1(\pi_p^*), \overline{u}_2(\pi_p^*)) = \left(p + \overline{u}_2(\pi_{p-1}^*), \frac{p+1}{p} \overline{u}_1(\pi_{p-1}^*) \right)$$

For any $\pi \in \mathcal{P}_p$, $p \geq 2$ we obtain a similar recursion that depends on whether π follows or deviates from π^* in extension of $\tilde{\pi}$ at each step. In the first case, the prefix of π is (12) and in the second case the prefix is (11). So we have

$$(\overline{u}_1(\pi), \overline{u}_2(\pi)) = \begin{cases} \left(p + \overline{u}_2(\tilde{\pi}), \frac{p+1}{p} \overline{u}_1(\tilde{\pi}) \right) & \text{if } \pi = 12\dots, \\ \left(p + \overline{u}_1(\tilde{\pi}), \frac{p+1}{p} \overline{u}_2(\tilde{\pi}) \right) & \text{if } \pi = 11\dots \end{cases}$$

$$\text{Then } (\overline{u}_1(\pi) - \overline{u}_1(\pi_p^*), \overline{u}_2(\pi) - \overline{u}_2(\pi_p^*)) =$$

$$\begin{cases} \left(\overline{u}_2(\tilde{\pi}) - \overline{u}_2(\pi_{p-1}^*), \frac{p+1}{p} (\overline{u}_1(\tilde{\pi}) - \overline{u}_1(\pi_{p-1}^*)) \right) & \text{if } \pi = 12\dots, \\ \left(\overline{u}_1(\tilde{\pi}) - \overline{u}_1(\pi_{p-1}^*), \frac{p+1}{p} (\overline{u}_2(\tilde{\pi}) - \overline{u}_2(\pi_{p-1}^*)) \right) & \text{if } \pi = 11\dots \end{cases}$$

We introduce notations to simplify the explanation. Using Proposition 1 we define $\delta_p = \overline{u}_1(\pi_p^*) - \overline{u}_2(\pi_p^*) = \frac{1}{3} [p + 1 + (-1)^{p+1} \gamma_{p+1}]$. We define the sets

$$A_p = \{ (\overline{u}_1(\pi) - \overline{u}_1(\pi_p^*), \overline{u}_2(\pi) - \overline{u}_2(\pi_p^*)) : \pi \in \mathcal{P}_p \}.$$

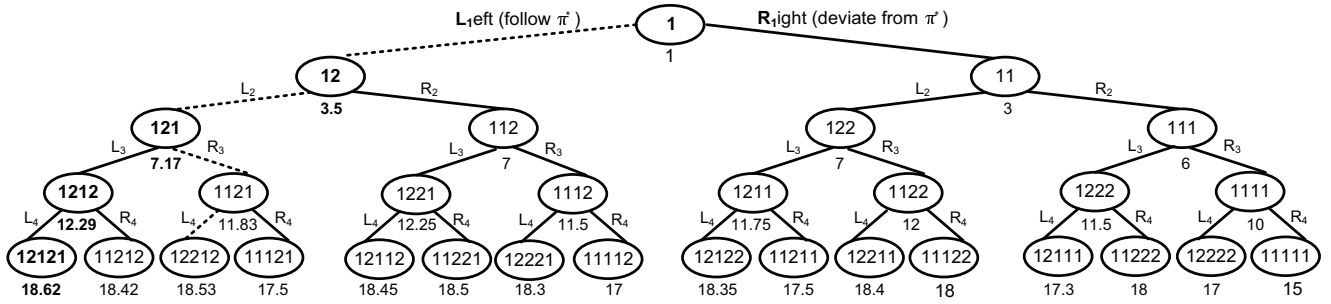


Figure 1: The policy tree of depth 5.

Note that for an element $(a, b) \in A_p$ corresponding to a policy $\pi \in \mathcal{P}_p$ we have $a + b = \overline{\text{sw}}(\pi) - \overline{\text{sw}}(\pi_p^*) = \text{dsw}_\pi$. Hence, π has a higher expected utilitarian social welfare than π_p^* if and only if $a + b > 0$.

The recursions above provide a description of the sets A_p . We have $A_1 = \{(0, 0)\}$ because π_1^* is the only policy of length 1, and for $p \geq 2$ the set A_p consists of the elements $(b, \frac{p+1}{p}a)$ and $(a + \delta_{p-1}, \frac{p+1}{p}(b - \delta_{p-1}))$ where (a, b) runs over A_{p-1} . Theorem 1 is equivalent to the following statement.

Proposition 6. Let $A_1 = \{(0, 0)\}$ and

$$A_k = \left\{ \left(b, \frac{k+1}{k}a \right) \right\} \cup \left\{ \left(a + \delta_k, \frac{k+1}{k}(b - \delta_k) \right) \right\}$$

for $k \geq 2$ where $(a, b) \in A_{k-1}$, $\delta_k = \frac{1}{3}(k + (-1)^k \gamma_k)$. Then $a + b \leq 0$ for all $(a, b) \in \bigcup_k A_k$.

Figure 2 shows the sets A_k , $k = 1, \dots, 4$ in the policy tree.




Figure 2: The sets A_k , $k = 1, \dots, 4$ in the policy tree.

Proving optimality. We might try to prove Proposition 6 inductively by deriving $a + b \leq 0$ for the point $(a, b) \in A_p$ corresponding to policy π from $a' + b' \leq 0$ for $(a', b') \in A_{p-1}$ corresponding to policy $\tilde{\pi}$. Unfortunately, the induction hypothesis is too weak as the following example shows.

Example 5. Assume $(a', b') = (-12, 11.9) \in A_{10}$ corresponding to some policy $\pi \in \mathcal{P}_{10}$. Let $(a, b) \in A_{11}$ be obtained from $(-12, 11.9)$ by deviating from π_{11}^* . With $\delta_{11} = 2.7643$ we obtain $a + b = -9.2357 + 9.9662 > 0$. Thus (a', b') satisfies Proposition 6 while (a, b) violates it.

To remedy this problem we would like to strengthen the proposition, for example by proving $a + b \leq f(a, b)$ for all $(a, b) \in \bigcup_k A_k$ where f is some function with $f(x, y) \leq 0$ for all (x, y) . The difficulty of finding such a function is indicated by Figure 3 showing the set A_{10} . Different markers distinguish the points arising from A_9 by following π^* from those deviating from π^* .




Figure 3: The set A_{10} and a more detailed view of the region around the origin.

The key idea of our proof is to strengthen Proposition 6 in another direction. We describe this strengthening first and then outline the induction argument. The technical details of the proof are presented in the online Appendix D. Consider a policy π that is represented by a node n_π at level k in the policy tree. Instead of requiring the inequality $a + b \leq 0$ only for the point $(a, b) \in A_k$ that corresponds to policy π , we also require it for (i) all policies that lay on the path that follow only the right branches from n_π and (ii) all policies that lay on the path that starts from n_π by following the left branch once and then only follow the right branches. To formalize this idea, for $k \geq 1$ we define functions $f_k, g_k : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by $f_k(x, y) = (y, \frac{k+2}{k+1}x)$ and $g_k(x, y) = (x + \delta_{k+1}, \frac{k+2}{k+1}(y - \delta_{k+1}))$. Note that $A_{k+1} = f_k(A_k) \cup g_k(A_k)$ for all $k \geq 1$, as f_k encodes the case when we follow the left branch and g_k - the right branch. Figure 2 illustrates this correspondence. We also consider iterated compositions of these functions. For every $k \geq 1$ let G_{k0} denote the identity on \mathbb{R}^2 , i.e. $G_{k0}(x, y) = (x, y)$, and for $m \geq 1$ let G_{km} denote the function

$$G_{km} = g_{k+m-1} \circ g_{k+m-2} \circ \dots \circ g_k.$$

Applying G_{km} to the point $(a, b) \in A_k$ corresponding to $\pi \in \mathcal{P}_k$ gives the point $(a', b') \in A_{k+m}$ which corresponds

to the policy $\pi' \in \mathcal{P}_{k+m}$ that is obtained from π by following the right branch m times. For all $k \geq 1$ and $m \geq 1$, we define the function $F_{km} = G_{k+1, m-1} \circ f_k$. F_{km} corresponds to starting in level k , following the first left branch and then $m-1$ right branches. For $(x, y) \in A_k$, $G_{km}(x, y) \in A_{k+m}$ for $m \geq 0$ and $F_{km}(x, y) \in A_{k+m}$ for $m \geq 1$. Proposition 6 is a consequence of the following theorem.

Theorem 2. Let $A_1 = \{(0, 0)\}$ and

$$A_{k+1} = f_k(A_k) \cup g_k(A_k) = \left\{ \left(y, \frac{k+2}{k+1}x \right) \right\} \cup \left\{ \left(x + \delta_{k+1}, \frac{k+2}{k+1}(y - \delta_{k+1}) \right) \right\}$$

for $k \geq 1$ where $(x, y) \in A_k$, $\delta_k = \frac{1}{3}(k + (-1)^k \gamma_k)$. Then for every $k \geq 1$ and every $(x, y) \in A_k$ the following statements are true.

1. For all $m \geq 0$, if $(x', y') = G_{km}(x, y)$ then $x' + y' \leq 0$.
2. For all $m \geq 1$, if $(x', y') = F_{km}(x, y)$ then $x' + y' \leq 0$.




Figure 4: Schematic representation of the proof of Theorem 2.

Proof sketch. We provide the full proof in the Appendix C. We give a high-level overview. We start with a few technical lemmas to derive an explicit description of functions $(x', y') = F_{km}(x, y)$ and $(x'', y'') = G_{km}(x, y)$. This gives us explicit expressions for the sums $x' + y'$ and $x'' + y''$ in terms of x and y . Then we proceed through the induction proof. We summarize the induction step here. Suppose the statements of the theorem are already proved for all sets A_l with $l < k$. Let (x, y) be an arbitrary element of A_k . Suppose $(x, y) = f_{k-1}(\tilde{x}, \tilde{y})$ for some $(\tilde{x}, \tilde{y}) \in A_{k-1}$ (the case $(x, y) = g_{k-1}(\tilde{x}, \tilde{y})$ is similar). Figure 4 shows (\tilde{x}, \tilde{y}) and (x, y) that is obtained from (\tilde{x}, \tilde{y}) by following the left branch. By the induction hypothesis, $x' + y' \leq 0$ whenever $(x', y') \in \{G_{k-1, m}(\tilde{x}, \tilde{y}), F_{k-1, m}(\tilde{x}, \tilde{y})\}$. The corresponding nodes are highlighted in gray in Figure 4. To complete the induction step we need to show $x' + y' \leq 0$ for $(x', y') \in \{G_{km}(x, y), F_{km}(x, y)\}$. The corresponding nodes are indicated by dashed circles. The result for $(x', y') = G_{km}(x, y)$ (gray and dashed) follows immediately as $(x', y') = G_{km}(x, y) = F_{k-1, m+1}(\tilde{x}, \tilde{y})$. For $(x', y') = F_{km}(x, y)$ we first express $x' + y'$ in terms of \tilde{x} and \tilde{y} . Then, by induction $x'' + y'' \leq 0$ for $(x'', y'') = G_{k-1, m-1}(\tilde{x}, \tilde{y})$.

Inverting the representation of $x'' + y''$ in terms of \tilde{x} and \tilde{y} we derive a bound $\tilde{x} + \tilde{y} \leq -c(m) \leq 0$, depending on m , and this stronger bound is used to prove $x' + y' \leq 0$ for $(x', y') = F_{km}(x, y)$. \square

The extension of Theorem 2 to n agents is not straightforward. Firstly, it requires deriving exact recursions for the expected utility for an arbitrary p . This is not trivial, as Proposition 1 only provides asymptotics. An easier extension might be to other utility functions. The alternating policy is not optimal for all scoring functions. For example, it is not optimal for the k -approval scoring function which has $g(i) = 1$ for $i \leq k$ and 0 otherwise. However, we conjecture that ALTPOLICY is optimal for all convex scoring functions (which includes lexicographical scoring).

6 Strategic Behaviour

So far, we have supposed agents sincerely pick the most valuable item left. However, agents can sometimes improve their utility by picking less valuable items. To understand such strategic behaviour, we view this as a finite repeated game with perfect information. [Kohler and Chandrasekaran, 1971] proves that we can compute the subgame perfect Nash equilibrium for the alternating policy with two agents by simply reversing the policy and the preferences and playing the game backwards. More recently, [Kalinowski *et al.*, 2012] prove this holds for any policy with two agents.

We will exploit such reversal symmetry. We say that a policy π is *reversal symmetric* if and only the reversal of π , after interchanging the agents if necessary, equals π . The policies 1212 and 1221 are reversal symmetric, but 1121 is not. The next result follows quickly by expanding and rearranging expressions for the expected utilitarian social welfare using the fact that we can compute strategic play by simply reversing the policy and profile and supposing truthful behaviour.

Theorem 3. For two agents and any utility function, any reversal symmetric policy that maximizes the expected utilitarian social welfare for truthful behaviour also maximizes the expected utilitarian social welfare for strategic behaviour.

As the alternating policy is reversal symmetric, it follows that the alternating policy is also optimal for strategic behaviour. Unfortunately, the generalisation of these results to more than two agents is complex. Indeed, for an unbounded number of agents, computing the subgame perfect Nash equilibrium becomes PSPACE-hard [Kalinowski *et al.*, 2013b].

7 Conclusions

Supposing additive utilities, and full independence between agents, we have shown that we can compute the expected utility of a sequential allocation procedure in polynomial time for any utility function. Using this result, we have proven that the expected utilitarian social welfare for Borda utilities is maximized by the alternating policy in which two agents pick items in a fixed order. We have argued that this mechanism remains optimal when agents behave strategically. There remain open several important questions. For example, is the alternating policy optimal for more than two agents? What happens with non-additive utilities?

References

- [Bouveret and Lang, 2011] 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, 2011.
- [Brams and Fishburn, 2000] S.J. Brams and P.C. Fishburn. Fair division of indivisible items between two people with identical preferences: Envy-freeness, pareto-optimality, and equity. *Social Choice and Welfare*, 17(2):247–267, 2000.
- [Brams and Kaplan, 2004] S.J. Brams and T.R. Kaplan. Dividing the indivisible. *Journal of Theoretical Politics*, 16(2):143–173, 2004.
- [Brams *et al.*, 2003] S.J. Brams, P.H. Edelman, and P.C. Fishburn. Fair division of indivisible items. *Theory and Decision*, 55(2):147–180, 2003.
- [Brams *et al.*, 2012] S. Brams, D. Kilgour, and C. Klamler. The undercut procedure: an algorithm for the envy-free division of indivisible items. *Social Choice and Welfare*, 39(2-3):615–631, 2012.
- [Herreiner and Puppe, 2002] D.K. Herreiner and C. Puppe. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19(2):415–430, 2002.
- [Kalinowski *et al.*, 2012] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia. Elicitation-free protocols for allocating indivisible goods. Fourth International Workshop on Computational Social Choice, 2012.
- [Kalinowski *et al.*, 2013a] T. Kalinowski, N. Narodytska, and T. Walsh. A social welfare optimal sequential allocation procedure. Technical report, CoRR archive, available at <http://arxiv.org/abs/1304.5892>, 2013.
- [Kalinowski *et al.*, 2013b] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia. Strategic behavior when allocating indivisible goods sequentially. In *Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI-13)*, 2013.
- [Kohler and Chandrasekaran, 1971] D.A. Kohler and R. Chandrasekaran. A class of sequential games. *Operations Research*, 19(2):270–277, 1971.