

Complexity of Manipulating Sequential Allocation

Haris Aziz

Data61, CSIRO and UNSW
Sydney, Australia
haris.aziz@data61.csiro.au

Sylvain Bouveret

LIG - Grenoble INP
France
sylvain.bouveret@imag.fr

Jérôme Lang

Université Paris-Dauphine, PSL Research University
CNRS, LAMSADE, Paris, France
lang@lamsade.dauphine.fr

Simon Mackenzie

Carnegie Mellon University,
Pittsburg, USA
simonm@andrew.cmu.edu

Abstract

Sequential allocation is a simple allocation mechanism in which agents are given pre-specified turns in which they take one item among those that are still available. It has long been known that sequential allocation is not strategyproof. This raises the question of the complexity of computing a preference report that yields a higher utility than the truthful preference. We show that the problem is NP-complete for one manipulating agent with additive utilities and several non-manipulating agents. In doing so, we correct a wrong claim made in a previous paper. We then give two additional results. First, we present a polynomial-time algorithm for optimal manipulation when the manipulator has additive binary utilities. Second, we consider a stronger notion of manipulation whereby the untruthful outcome yields more utility than the truthful outcome for all utilities consistent with the ordinal preferences; for this notion, we show that a manipulation, if any, can be computed in polynomial time.

Introduction

A simple but popular mechanism to allocate indivisible items is *sequential allocation* (Aziz, Walsh, and Xia 2015; Bouveret and Lang 2011; Brams and Straffin 1979; Brams and Taylor 1996; Kalinowski, Narodytska, and Walsh 2013; Kalinowski et al. 2013; Kohler and Chandrasekaran 1971; Levine and Stange 2012; Tominaga, Todo, and Yokoo 2016). In sequential allocation, a sequence specifies the turns of the agents. For example, for sequence 1212, agents 1 and 2 alternate with agent 1 taking the first turn. Agents report their preferences by expressing a linear order over items, and are allocated items based on their reported preferences as follows. They get turns according to the sequence, and when their turn comes, they are given the most preferred item (according to their reported order) that has not yet been allocated. Sequential allocation is an ordinal mechanism since the outcome only depends on the ordinal preferences of agents over items. Nevertheless, it is a standard assumption

in the literature that agents have underlying additive utilities for the items.¹ It has long been known that sequential allocation is not strategyproof in particular when agents do not have consecutive turns. This motivates the natural problem of computing best responses (also referred to as manipulations). Kohler and Chandrasekaran (1971) presented a polynomial-time algorithm to compute the optimal manipulation of an agent when there are two agents and the sequence is alternating (121212..). Bouveret and Lang (2011) initiated further work on manipulation of sequential allocation and showed that (1) it can be checked in polynomial time whether an agent can be allocated a given subset of items; then, they show in a later work (Bouveret and Lang 2014) that (2) an optimal manipulation for two agents (one manipulator and one non-manipulator) can be found in polynomial time for any sequence. Now, they also claim (page 142, left column, lines 5-7) that by putting together (1) and (2), a best response can be computed in polynomial time for one manipulator against *several* non-manipulators, and that each best response results in the same allocation for the manipulator. As we will show, this conclusion is wrong.

Results We focus on computing best responses (or manipulations) under sequential allocation. We first show that the algorithm given by Bouveret and Lang (2014) for computing a best response against one non-manipulator does not extend to several non-manipulators. We then show that the problem of computing a best response is in fact NP-complete. The result has some interesting consequences since many allocation rules are based on sequential allocation. We also give two additional results. We first present a polynomial-time algorithm for optimal manipulation when the manipulator has binary utilities. We then consider a stronger notion of manip-

¹We view here sequential allocation as a centralized mechanism; it can also be seen as a decentralized mechanism, where instead of reporting their preference, agents pick a remaining item each time their turn come. Our results hold for both views but are easier to state for the centralized view.

ulation whereby the untruthful outcome yields more utility than the truthful outcome for all utilities consistent with the ordinal preferences. For this notion, we give a polynomial-time algorithm for computing a manipulation.

Preliminaries

We consider a set of agents $N = \{1, \dots, n\}$ and a set of items $O = \{o_1, \dots, o_m\}$. Agent 1 is the *manipulator* and agents 2 to n the *non-manipulators*. The preference profile $\succ = (\succ_1, \dots, \succ_n)$ specifies for each agent i his complete, strict, and transitive preference \succ_i over O . A preference relation \succ_i is usually denoted by a list, such as $i : a, b, c, d$ for $a \succ_i b \succ_i c \succ_i d$. Moreover, for agent 1 we are also given a set of positive values $\{u_1(o) \mid o \in O\}$, where for each o , $u_1(o) > 0$ is 1's cardinal evaluation of the value of item o . These values are consistent with \succ_1 , that is, $u_1(o) > u_1(o')$ iff $o \succ_1 o'$. We assume that agent 1's preference over sets of items are represented by the additive utility function $u_i(O') = \sum_{o \in O'} u_i(o)$ for all $O' \subseteq O$. We say that u_1 is *lexicographic* if for all $o \in O$, $u_1(o) > \sum_{o' \succ_i o} u_1(o')$. We abbreviate $u_1(S) \geq u_1(T)$ into $S \succeq_1 T$. Finally, an assignment p is a function from O to N .

Complexity of Manipulation

We first show that the best response algorithm of Bouveret and Lang (2014) does not work for $n \geq 3$, and that several optimal manipulations may result in different allocations.

Example 1. *Let the sequence be 1231, and consider the following profile \succ .*

$$\succ_1: a, b, c, d \quad \succ_2: c, d, a, b \quad \succ_3: a, b, c, d$$

According to (Bouveret and Lang 2014), the best response is one in which agent 1 gets $\{a, d\}$ which can even be achieved by the truthful report. Now, if agent 1 misreports

$$\succ'_1: c, b, a, d$$

then he gets $\{b, c\}$. Which one of $\{a, d\}$ and $\{b, c\}$ is better depends on u_1 and not just on \succ_1 . For example, if $u_1(a) = 5$, $u_1(b) = 4$, $u_1(c) = 3$ and $u_1(d) = 1$, then $\{b, c\} \succ_1 \{a, d\}$ and the best achievable allocation for 1 is $\{b, c\}$. On the other hand, if $u_1(a) = 4$, $u_1(b) = 3$, $u_1(c) = 2$ and $u_1(d) = 1$, then 1 is indifferent between $\{a, d\}$ and $\{b, c\}$.

Thus, for more than one non-manipulator the algorithm of Bouveret and Lang (2014) does not necessarily compute a best response,² and the complexity of computing a best response is an open problem. We are now going to show that the problem is NP-hard. The reduction involves a similar high-level idea as that of the result of Aziz et al. (2015a) that manipulating the probabilistic serial (PS) mechanism is NP-hard. However, the reduction requires new gadgets. Also note that the NP-hardness result for the PS mechanism does

²The reason is that when they translate the n -agent problem P of deciding whether agent 1 can obtain a set of objects S into a 2-agent problem P' (Bouveret and Lang 2011), the preference relation of agent 2 in P' problem depends on S , and in the best response algorithm for a 2-agent problem of Bouveret and Lang (2014), S is not fixed.

not directly imply a similar result for sequential allocation. Similarly, the NP-hardness of the manipulation of sequential allocation does not imply the NP-hardness of the manipulation of the PS mechanism.

Theorem 2. *Computing a best response for the sequential allocation mechanism is NP-complete.*

Proof. We prove that the following problem (BEST RESPONSE) is NP-complete: given an assignment setting, a utility function $u_1 : O \rightarrow \mathbb{N}$ for the manipulator, and a target utility T , can the manipulator specify preferences such that the utility for his allocation under the sequential allocation rule is at least T ? The problem BEST RESPONSE is clearly in NP, since the outcome with respect to the reported preference can be computed by simulating sequential allocation.

We show hardness by reduction from a restricted NP-complete version of 3SAT where each literal appears exactly twice in the set of clauses. The problem remains NP-complete (Berman, Karpinski, and Scott 2003). Given such a 3SAT instance $F = (X, C)$ where $X = \{x_1, \dots, x_{|X|}\}$ is the set of variables and C the set of clauses, we build an instance of BEST RESPONSE where the manipulator can obtain utility $\geq T$ if and only if C is satisfiable. We denote by $L = \{x_1, \neg x_1, \dots, x_{|X|}, \neg x_{|X|}\}$ the set of literals.

We define the set of agents as

$$N = \{1\} \cup \{a_{l_i}^1, a_{l_i}^2 : l_i \in L\}$$

with agent 1 as the manipulator, and the set of items as

$$O = \{o_c^1, o_c^2, o_c^3 : c \in C\} \\ \cup \{o_{l_i}^1, o_{l_i}^2, h_{l_i}^1, h_{l_i}^2, h_{l_i}^3, d_{l_i}^{11}, d_{l_i}^{12}, d_{l_i}^{21}, d_{l_i}^{22} : l_i \in L\}$$

with

- $O(c) = \{o_c^1, o_c^2, o_c^3\}$: clause items associated with clause c
- $O(l_i) = \{o_{l_i}^1, o_{l_i}^2\}$: choice items associated with literal l_i
- $H(l_i) = \{h_{l_i}^1, h_{l_i}^2, h_{l_i}^3\}$: consistency items associated with l_i
- $D(l_i) = \{d_{l_i}^{11}, d_{l_i}^{12}, d_{l_i}^{21}, d_{l_i}^{22}\}$: dummy items associated with l_i

We view the sequential allocation process as follows. The agents' preferences are built in a way such that the agents will first go through $|X|$ choice rounds corresponding to variables $x_1, \dots, x_{|X|}$, then $|C|$ clause rounds corresponding to $c_1, \dots, c_{|C|}$, and then one final collection round.

High-level Idea The picking sequence is composed of successive rounds. Each round is associated with a subset of items that are valued by agent 1 much higher than items associated with later rounds. If he wants to reach the target utility T , then in each round, agent 1 must focus only on the items relevant to that round; those he will not get during the corresponding round will be taken by other agents and it will not be possible for him to take them in later rounds. In each round, agent 1 makes a choice between the items corresponding to a variable and to its negation. There is a

negligible difference between the utility of the items corresponding to the variable and those corresponding its negation (for example o_x^1 and $o_{\neg x}^1$). If he chooses the items in a ‘correct’ way, then he will get a most preferred item corresponding to each of the clauses: more precisely, agent 1, in any case, will get one item $h_{x_i}^j$ and one item $h_{\neg x_i}^k$ for each variable x_i ; if he chooses in a ‘correct’ way, then this pair of items $\{h_{x_i}^j, h_{\neg x_i}^k\}$ will give him a sufficiently high utility, otherwise there will be a utility loss which he will never be able to compensate. We show below that agent 1 reaches utility T if and only if these two conditions are satisfied:

1. he chooses the choice items in a consistent way, that is, at each choice round i he chooses either $o_{x_i}^1$ and $o_{x_i}^2$ (which we view as assigning x_i to false), or $o_{\neg x_i}^1$ and $o_{\neg x_i}^2$ (which we view as assigning x_i to true).
2. he manages to get his most preferred clause item for each clause c ; this is possible only if none of the agents corresponding to the negation of the literals in c gets this clause item before him.

Utility function of agent 1 Agent 1 highly prefers choice items relevant to variable x_i to items relevant to variable x_j , $j > i$; for a given variable he highly prefers its relevant choice items to its relevant consistency items. These items are highly preferred to clause items, which in turn are highly preferred to dummy items:

$$\begin{aligned} O(x_1), O(\neg x_1) &\gg_1 H(x_1), H(\neg x_1) \\ &\gg_1 \dots \gg_1 O(x_{|X|}), O(\neg x_{|X|}) \gg_1 H(x_{|X|}), H(\neg x_{|X|}) \\ &\gg_1 o_{c_1}^1 \gg_1 o_{c_2}^1 \gg_1 \dots \gg_1 o_{c_{|C|}}^1 \gg_1 \text{other items} \end{aligned}$$

His utilities for items in $O(x_i)$ and $H(x_i)$ are as follows:

- $u_1(o_{x_i}^1) = u_1(o_{\neg x_i}^1) + \epsilon \gg u_1(o_{x_i}^2) = u_1(o_{\neg x_i}^2) + \epsilon$, where ϵ is a negligible quantity.
- $h_{\neg x_i}^1 \succ_1 h_{\neg x_i}^2 \succ_1 h_{\neg x_i}^3 \succ_1 h_{x_i}^1 \succ_1 h_{x_i}^2 \succ_1 h_{x_i}^3$.
- $u_1(h_{x_i}^2) + u_1(h_{\neg x_i}^2) < u_1(h_{\neg x_i}^1) + u_1(h_{x_i}^1) = u_1(h_{x_i}^1) + u_1(h_{\neg x_i}^1)$.
- agent 1 reaches utility T if and only if he gets two items corresponding to a variable, at least one top choice consistency item (h_x^1 or $h_{\neg x}^1$) in each round and his target clause item in each clause round.

Choice Round The choice round is composed of a series of rounds, one for each variable x_i ; each of them is itself composed of two subrounds; in each of them, agent 1 has to choose a choice item. The sub-sequence in choice round corresponding to variable x_i is as follows:

$$\begin{aligned} &1, a_{\neg x_i}^1, a_{\neg x_i}^2, a_{x_i}^1, a_{x_i}^2, 1, a_{\neg x_i}^1, a_{\neg x_i}^2, a_{x_i}^1, \\ &a_{x_i}^2, a_{\neg x_i}^1, a_{\neg x_i}^2, 1, a_{x_i}^1, a_{x_i}^2, 1 \end{aligned}$$

The ordinal preferences relevant for the choice round corresponding to variable x_i are as follows. Recall that agent 1’s preferences are

$$1 : o_{x_i}^1, o_{\neg x_i}^1, o_{x_i}^2, o_{\neg x_i}^2, h_{\neg x_i}^1, h_{\neg x_i}^2, h_{\neg x_i}^3, h_{x_i}^1, h_{x_i}^2, h_{x_i}^3$$

Then, each literal agent prefers the items corresponding to the negation of the literal:

$$\begin{array}{l} a_{\neg x_i}^1 : o_{x_i}^1, d_{x_i}^{11}, d_{x_i}^{12}, o_{x_i}^2, h_{x_i}^1, h_{x_i}^2, h_{x_i}^3 \\ a_{x_i}^2 : d_{x_i}^{21}, o_{x_i}^1, o_{x_i}^2, d_{x_i}^{22}, h_{x_i}^1, h_{x_i}^2, h_{x_i}^3 \\ \hline a_{x_i}^1 : o_{\neg x_i}^1, d_{\neg x_i}^{11}, h_{\neg x_i}^1, o_{\neg x_i}^2, h_{\neg x_i}^2, h_{\neg x_i}^3, d_{\neg x_i}^{12} \\ a_{x_i}^2 : d_{\neg x_i}^{21}, o_{\neg x_i}^1, o_{\neg x_i}^2, h_{\neg x_i}^1, h_{\neg x_i}^2, h_{\neg x_i}^3, d_{\neg x_i}^{22} \end{array}$$

Items that do not appear in these lists are below in the preference list of an agent. Note the asymmetry between agents corresponding to positive and negative literals: the positive (respectively, negative) literal agents have a dummy item as the least preferred item (respectively, the consistency items as the least preferred items) relevant to the picking in the choice round.

We say that agent 1 makes a *consistent choice* in the choice round corresponding to x_i if he picks both $o_{\neg x_i}^1$ and $o_{x_i}^2$ (in which case we say he *assigns* x_i to true) or both $o_{x_i}^1$ and $o_{\neg x_i}^2$ (in which case we say he *assigns* x_i to false). Given the preferences of other agents, the pairs that agent 1 can possibly get are only these three ones: $\{h_{x_i}^1, h_{\neg x_i}^3\}$, $\{h_{\neg x_i}^1, h_{x_i}^3\}$ and $\{h_{x_i}^2, h_{\neg x_i}^2\}$. If he makes a consistent choice then he will get one of the first two pairs, for which he has the same utility (see Tables 1 and 2), and if he makes an inconsistent choice then he will get the pair $\{h_{x_i}^2, h_{\neg x_i}^2\}$, which gives him a smaller utility than the other two pairs (see Tables 3 and 4).

Clause round The sequence in clause round corresponding to clause $c = (l_i \vee l_j \vee l_k)$ is as follows.

$$a_{l_i}, a_{l_j}, a_{l_k}, - - -, 1$$

For each literal l in clause c , there is an agent $a_{\neg l}^1$ or $a_{\neg l}^2$ that features in the round. Recall that each literal occurs in exactly two clauses in the set of clauses. Agent $a_{\neg l}^1$ (respectively, $a_{\neg l}^2$) features if c is the first (respectively, second) clause in which l is present.

After all the clause rounds are finished, agent 1 will get $|C|$ turns in which he will get the clause items that are still available.

For agent 1, according to u_1 , his relevant preference in the clause round is to go for o_c^1 .

For a literal l , let c and c' be the two clauses containing l . If $l = x$ (x being a propositional variable), the preferences of agents $a_{\neg x}^1$ and $a_{\neg x}^2$ in the round corresponding to l are

$$\begin{array}{l} a_{\neg x}^1 : h_x^2, h_x^3, o_c^3, o_c^2, o_c^1 \\ a_{\neg x}^2 : h_x^2, h_x^3, o_{c'}^3, o_{c'}^2, o_{c'}^1 \end{array}$$

and if $l = \neg x$ then the preferences of agents a_x^1 and a_x^2 in the round corresponding to l are

$$\begin{array}{l} a_x^1 : d_{\neg x}^{12}, o_c^3, o_c^2, o_c^1 \\ a_x^2 : d_{\neg x}^{22}, o_{c'}^3, o_{c'}^2, o_{c'}^1 \end{array}$$

Note that the consistency items h_x^2 and h_x^3 , as well as the dummy items $d_{\neg x}^{12}$ and $d_{\neg x}^{22}$, were possibly picked in the choice round corresponding to variable x .

Stage	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Agent	1	$a_{x_i}^1$	$a_{\neg x_i}^2$	$a_{x_i}^1$	$a_{\neg x_i}^2$	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	1	$a_{x_i}^1$	$a_{\neg x_i}^2$	1
Item picked	$o_{\neg x_i}^1$	$o_{x_i}^2$	$d_{x_i}^{21}$	$d_{\neg x_i}^{11}$	$d_{\neg x_i}^{21}$	$o_{\neg x_i}^2$	$d_{x_i}^{11}$	$o_{x_i}^2$	$h_{\neg x_i}^1$	$h_{x_i}^2$	$d_{x_i}^{12}$	$d_{x_i}^{22}$	$h_{\neg x_i}^3$	$d_{\neg x_i}^{12}$	$d_{\neg x_i}^{22}$	$h_{x_i}^1$

Table 1: Choice round for variable x_i in which agent 1 makes consistent choice $\{o_{\neg x_i}^1, o_{x_i}^2\}$ (“assign x_i to true”). Agents $a_{\neg x_i}^1$ and $a_{x_i}^2$ next focus on items $h_{x_i}^2$ and $h_{x_i}^3$ before turning their attention to the clause items. On the other hand, $a_{x_i}^1$ and $a_{\neg x_i}^2$ are now ready to get clause items. In this way, the agents corresponding to literals set to false are quicker to get their clause items.

Stage	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Agent	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	1
Item picked	$o_{x_i}^1$	$d_{x_i}^{11}$	$d_{x_i}^{21}$	$o_{\neg x_i}^1$	$d_{\neg x_i}^{21}$	$o_{x_i}^2$	$d_{x_i}^{12}$	$d_{x_i}^{22}$	$d_{\neg x_i}^{11}$	$o_{\neg x_i}^2$	$h_{x_i}^1$	$h_{x_i}^2$	$h_{\neg x_i}^1$	$h_{\neg x_i}^2$	$h_{x_i}^3$	$h_{x_i}^1$

Table 2: Choice round for variable x_i in which agent 1 makes consistent choice $\{o_{x_i}^1, o_{\neg x_i}^2\}$ (“assign x_i to false”). Agents $a_{x_i}^1$ and $a_{\neg x_i}^2$ next focus on items $d_{\neg x_i}^{12}$ and $d_{\neg x_i}^{22}$ before turning their attention to the clause items. On the other hand, $a_{\neg x_i}^1$ and $a_{x_i}^2$ are now ready to get clause items.

Stage	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Agent	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	1
Item picked	$o_{x_i}^1$	$d_{x_i}^{11}$	$d_{x_i}^{21}$	$o_{\neg x_i}^1$	$d_{\neg x_i}^{21}$	$o_{\neg x_i}^2$	$d_{x_i}^{12}$	$o_{x_i}^2$	$d_{\neg x_i}^{11}$	$h_{\neg x_i}^1$	$h_{x_i}^1$	$d_{x_i}^{22}$	$h_{\neg x_i}^2$	$h_{\neg x_i}^3$	$h_{x_i}^3$	$h_{x_i}^2$

Table 3: Choice round for variable x_i in which agent 1 makes inconsistent choice $\{o_{x_i}^1, o_{\neg x_i}^2\}$. As a result of making an inconsistent choice, agent 1 does not get $h_{\neg x_i}^1$ nor $h_{x_i}^1$.

Stage	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Agent	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	$a_{\neg x_i}^1$	$a_{x_i}^2$	1	$a_{\neg x_i}^1$	$a_{x_i}^2$	1
Item picked	$o_{\neg x_i}^1$	$o_{x_i}^2$	$d_{x_i}^{21}$	$d_{\neg x_i}^{11}$	$d_{\neg x_i}^{21}$	$o_{x_i}^2$	$d_{x_i}^{11}$	$d_{x_i}^{22}$	$h_{\neg x_i}^1$	$o_{\neg x_i}^2$	$d_{x_i}^{12}$	$h_{x_i}^1$	$h_{\neg x_i}^2$	$d_{\neg x_i}^{12}$	$d_{\neg x_i}^{22}$	$h_{x_i}^2$

Table 4: Choice round for variable x_i in which agent 1 makes inconsistent choice $\{o_{\neg x_i}^1, o_{x_i}^2\}$.

If agent 1 has made a consistent choice in the choice round corresponding to variable x and has set x to “true”, then $a_{x_i}^1$ and $a_{\neg x_i}^2$ are now ready to get clause items in their respective clause rounds. If agent 1 has made a consistent choice in the choice round corresponding to variable x and has set x to “false”, then $a_{\neg x_i}^1$ and $a_{x_i}^2$ already want to get clause items in their respective clause rounds.

In the clause round, for any literal l that is not satisfied by the truth assignment resulting from the choice round, the agent corresponding to l gets a clause item o_c^i where c is a clause containing $\neg l$. For example, if x has been assigned to false, then $a_{\neg x}^1$ gets a clause item for the first clause in which x is present and $a_{x_i}^2$ gets a clause item for the second clause in which x is present. Therefore, if all the literals of a clause c are false, then agent 1 does not get o_c^1 .

If literal l is satisfied, the agent $a_{\neg l}$ corresponding to it gets a dummy or consistency item in that round instead of a clause item. This means that if all the literals of a clause are not satisfied, then all three clause items of a clause are gone, and agent 1 does not get a clause item. He is much more interested in the clause item o_c^1 than in the clause items o_c^2 and o_c^3 . The other clause items are far down in his preference list so he would rather get all the top clause items o_c^1 for each clause c rather than o_c^2 and o_c^3 .

For example, let us consider clause $c = (x_i \vee \neg x_j \vee \neg x_k)$, and assume variables x_i , x_j and x_k are set to true, *i.e.*, agent 1 has picked choice items corresponding to $\neg x_i$, $\neg x_j$ and $\neg x_k$. This means that in the clause round, $a_{x_j}^1$ and $a_{x_k}^1$ are ready to take their clause items o_c^3 and o_c^2 but $a_{\neg x_i}^1$ wants to get one of the unallocated consistency items before he is interested in consistency items o_c^3 , o_c^2 , o_c^1 . This is helpful for agent 1 because he can get o_c^1 . Since each literal occurs

exactly twice in the set of clauses, note that as long as 1 makes a consistent choice, there will be another clause c' in which literal x is present and if x is set to true, then $a_{\neg x}^2$ will get $h_{x_i}^3$ and hence 1 will be able to get $o_{c'}$. Note that after the clause rounds all the consistency items are already allocated, so agent 1 can hope to get all the top clause items if they were not already taken in the clause rounds.

Stage	1	2	3	after clause rounds
Agent	$a_{\neg x_i}^1$	$a_{x_j}^1$	$a_{x_k}^1$	1
Item picked	$h_{x_i}^2$	o_c^2	o_c^3	o_c^1

Table 5: Clause round corresponding to $c = (x_i \vee \neg x_j \vee \neg x_k)$

Collection round The sequence in this round is $1, \dots, 1$.

At this stage, agent 1 prefers $o_{c_1}^1, \dots, o_{c_{|C|}}^1$ to all other remaining items.

The idea is that if agent 1 make choices that sets all the clauses to true, then he gets all the clause items. Note that if 1 makes a consistent choice for the variables but does not pick up all the clause items in the collection round (because the set of clauses is unsatisfiable), then 1 does not get all the clause items. Since there are items less preferred by 1 than the o_c^1 s such as o_c^2 and o_c^3 s, agent 1 is forced to pick a much less preferred item in the collection round.

Based on construction of the choice and clause rounds, we are in a position to prove a series of claims.

Claim 1. *If agent 1 does not make a consistent choice of the variable items, then he does not get utility T .*

Proof. If 1 does not make a choice at all in a choice round (that is, if he does not pick one of o_x^1 or $o_{\neg x}^1$ and one of o_x^2 or $o_{\neg x}^2$), then his most preferred items corresponding to the literals are then taken by the agents corresponding to the literal. If 1 makes a choice in each choice round but not a consistent one, then he gets $\{h_{\neg x}^2, h_x^2\}$ which has sufficiently less utility than $\{h_{\neg x}^1, h_x^3\}$ or $\{h_x^1, h_{\neg x}^3\}$ which implies he cannot get total utility T . \square

Claim 2. *If agent 1 makes consistent choices but the assignment is not satisfying, then agent 1 does not get utility T .*

Proof. If some clause c is set to false, then agent 1 is not able to o_c^1 because the literal agents in the clause round corresponding to c take all the items o_c^3, o_c^2, o_c^1 . This implies that agent 1 does not get utility T . \square

Claim 3. *If there exists a satisfying assignment for C , then agent 1 can get utility T .*

Proof. If there exists a satisfying assignment, then consider the preference report of agent 1 in which in each choice round, he picks $o_{x_i}^1$ and $o_{x_i}^2$ if x_i is set to be false. By doing this he gets to pick a top consistency item in that round as well. Since all the clauses are satisfied, in each clause round, agent 1 is able to get his clause item o_c^1 . The utilities are set in a way so that as long as agent 1 gets two items corresponding to the same literal and hence at least one top choice consistency items in each round and his target clause item in each clause round, agent 1 gets utility at least T . \square

The claims show that agent 1 gets utility at least T if and only if there is a satisfying truth assignment. \square

The result above also gives us the following statement: computing a best response is NP-hard. This raises the following question: what is the complexity of testing whether there exists a report that gives more utility than the truthful report?

Theorem 3. *Checking whether there exists a report that yields more utility than the truthful report is NP-complete.*

Proof. We first prove that the restriction of SAT to set of clauses such that (a) exactly one clause has only negative literals, and (b) each literal appears exactly twice in the set of clauses, is NP-complete. We show this by reduction from SAT with restriction (b), which we know to be NP-complete. Let $F = \{c_1, \dots, c_n\}$ be the initial set of clauses, with (b) holding. Let G the following set of clauses, where d_1, \dots, d_n are new variables:

$$G = \{c_i \vee d_i, 1 \leq i \leq n\} \cup \{\neg d_i \vee d_{i+1}, 1 \leq i < n\} \\ \cup \{\neg d_n\} \cup \{d_1 \vee \neg d_1 \vee \dots \vee \neg d_n\}$$

Each literal (including d_i and $\neg d_i$ for all i) appears exactly twice in G , all clauses except one ($\neg d_n$) contain at least one positive literal. Now, G is equivalent to $\{\neg d_1, \dots, \neg d_n\} \cup F$, therefore G is satisfiable if and only if F is satisfiable.

We now present a reduction from SAT with restrictions (a) and (b) to the existence of a better response than the truthful report. The reduction is essentially the same as that in

the proof of Theorem 2, except that the clauses may not be of size 3. For clauses of different size, we simply create as many clause items as the number of literals in the clause and agent 1 still needs to get the top clause item of each clause to achieve his target utility. The reduction again involves choice rounds for all the variables and the goal is for the agent 1 to maximize utility by making consistent choices and then get as many clause items as possible.

In the same way as in the proof of Theorem 2, we can show that if the set of clauses is not satisfiable, agent 1's optimal strategy is to tell the truth (*i.e.*, set all variables to true) because this leads him to get consistent variable items and all clause items except for the last clause, and that if the set of clauses is satisfiable, then the satisfying truth assignment also gives the corresponding optimal preference report for agent 1. Therefore, checking whether there exists a better report for agent 1 than the truthful report is NP-complete. \square

Manipulation under Binary Utilities

We now consider the restriction that the manipulator is asked to report a linear order but has *binary* preferences over items: the set of items is partitioned into two subsets O^+ and O^- , such that for each $o \in O^+$, $u_i(o) = \alpha$ and for each $o \in O^-$, $u_i(o) = \beta < \alpha$; that is, the items in O^+ are the most preferred items of agent 1. Binary utilities are important for *e.g.*, when each agent is only interested in a subset of acceptable items (Bogomolnaia and Moulin 2004). Since the number of items he gets with a fixed sequence is fixed, the utility obtained with a given report is maximal if and only if the number of objects in O^+ he gets is maximal. Therefore the manipulation problem under this restriction consists in finding a report leading agent 1 to get a maximum number of 'top items'.

Let us assume that agent 1 (the manipulator) has n_1 turns. Let π_{-1} the policy obtained by removing all occurrences of 1 in π , and let $First(\pi_{-1}, \succ_2, \dots, \succ_n, O^+)$ be the first item in O^+ picked by some agent when simulating π_{-1} on $(\succ_2, \dots, \succ_n)$. For example, take $\pi = 231312$, $\succ_2 = (o_1, o_2, o_3, o_4, o_5, o_6)$, $\succ_3 = (o_1, o_3, o_5, o_2, o_4, o_6)$ and $O^+ = \{o_2, o_4\}$. Then we have $\pi_{-1} = 2332$ and $First(\pi_{-1}, \succ_2, \succ_3, O^+) = o_2$, since simulating π_{-1} leads to 2 taking o_1 , 3 taking o_3 and o_5 , and then 2 taking o_2 .

Moreover, let us write $\pi = (head(\pi), 1, tail(\pi))$, where $head(\pi)$ is the longest starting subsequence of π in which 1 does not occur, and $tail(\pi)$ is the subsequence of π starting right after the first occurrence of 1. Let $Allocate(head(\pi), \succ_2, \dots, \succ_n)$ be the set of items allocated to 2, \dots , n when simulating $head(\pi)$ with $(\succ_2, \dots, \succ_n)$. For instance, with π , \succ_2 and \succ_3 as above, we have $head(\pi) = 23$, $tail(\pi) = 312$, and $Allocate(head(\pi), \succ_2, \succ_3) = \{o_1, o_3\}$.

Consider the following algorithm *BR*:

Input: O^+ , π , $(\succ_2, \dots, \succ_n)$

Output: \succ'_1

$k \leftarrow 1$;

$n_1 \leftarrow$ number of occurrences of 1 in π ;

While $k \leq n_1$ and $O^+ \neq \emptyset$ and $\pi \neq 11 \dots 1$

$O^* \leftarrow Allocate(head(\pi), \succ_2, \dots, \succ_n)$;

remove O^* from O, O^+ , and $(\succ_2, \dots, \succ_n)$;
 $\pi \leftarrow \text{tail}(\pi)$;
 $a_k \leftarrow \text{First}(\pi_{-1}, \succ_2, \dots, \succ_n, O^+)$;
remove a_k from O, O^+ , and $(\succ_2, \dots, \succ_n)$;
 $k \leftarrow k + 1$;

End While

$\succ'_1 \leftarrow (a_1, \dots, a_k)$;

complete \succ'_1 with the remaining items of O^+ (if any) in an arbitrary way, and then by other items in an arbitrary way;

Return \succ'_1

Theorem 4. *A best response can be computed in polynomial time when the manipulator has binary utilities: Algorithm BR outputs a best response.*

Proof. Assume that agent 1 has n_1 turns. 1 is interested in getting a maximum number n'_1 top items from O^+ . Let \succ'_1 be the sequence of objects obtained using Algorithm BR.

The main idea of the algorithm is that \succ'_1 is gradually built in a way so that 1 can get all the top items in the list. A new top item a_k is only appended to the list if 1 can additionally get it along with the previous items in the list. Just before adding a_k , when we simulate sequential allocation, 1 does not get a_k because someone else gets it right after 1's turn in which 1 abstained from picking an item because \succ'_1 only had $k-1$ items. Hence if 1 does not abstain and in fact makes use of his k -th turn, then 1 can get a_k .

To prove the optimality of the algorithm, we reason iteratively on the number of picking turns of agent 1. Namely, we prove the following statement: (H_{n_1}) suppose that the number of remaining picking turns of 1 is n_1 . Suppose that there is a picking strategy for agent 1 to get c objects among those from O^+ , and let $\succ''_1 = \langle b_1, \dots, b_{n_1} \rangle$ be the sequence of objects he gets with this strategy. Then the strategy $\succ'''_1 = \langle a_1, b_2, \dots, b_{n_1} \rangle$, where we have simply replaced b_1 by a_1 in \succ''_1 gives at least c objects from O^+ to agent 1.

We omit the argument of (H_{n_1}) due to lack of space.

By successively applying this hypothesis on $\langle b_{n_1} \rangle, \langle b_{n_1-1}, b_{n_1} \rangle, \dots$ we prove that for each k , the strategy $\langle a_k, \dots, a_{n_1} \rangle$ gives to agent 1 as many objects from O^+ as the strategy $\langle b_k, \dots, b_{n_1} \rangle$. This proves that for each arbitrary picking strategy, the strategy given by the algorithm gives to agent 1 at least as many objects from O^+ . Hence, the algorithm returns an optimal strategy. \square

Manipulation under responsive preferences

An allocation S is *more preferred with respect to responsive (RS) preferences* than allocation T if S is a result of replacing an item in T with a strictly more preferred item (Aziz et al. 2015b; Brams and King 2005, for instance). Note that the responsive relation is transitive but not complete.

We examine the problem of determining whether there exists an untruthful report which yields an allocation that is strictly more preferred with respect to responsive preferences. The problem is equivalent to checking whether there exists a report that yields more utility with respect to *all* additive utilities consistent with the ordinal preferences. In contrast to the problem for specific utilities, this particular problem of checking whether there exists such a clear manipulation can be solved in polynomial time. Our algorithm

is based on an intimate connection that we identify between manipulation under responsive preferences and under lexicographic preferences.

Theorem 5. *It can be checked in polynomial time whether there exists a manipulation that gives an allocation that is strictly more preferred with respect to responsive set extension than the truthful outcome.*

Proof. We first present the algorithm.

Input: picking sequence π , profile $(\succ_1, \succ_2, \dots, \succ_n)$

Output: \succ'_1

Rename items such that $\succ_1 = o_1, \dots, o_m$

$O_1 \leftarrow$ set of items obtained by agent 1 with his sincere strategy;

$k \leftarrow 1; O^* \leftarrow \emptyset$;

Repeat

If agent 1 can obtain $O^* \cup \{o_k\}$ **then**

$O^* \leftarrow O^* \cup \{o_k\}$

endif

$k \leftarrow k + 1$

Until O^* lexicographically dominates O_1 **or** $k > m$

If O^* lexicographically dominates O_1 **then**

$\succ'_1 \leftarrow$ ranking of O^* allowing agent 1 to obtain O^*

Complete \succ'_1 with the items in $O \setminus O^*$, ranked as in \succ_1

Return \succ'_1

else Return 'failure'

endif

It can be determined in polynomial time whether agent 1 can obtain $O^* \cup \{o_k\}$ (Bouveret and Lang 2011, Proposition 7), therefore the algorithm works in polynomial time. Now we claim that when the algorithm returns \succ'_1 , then \succ'_1 yields a strictly better allocation than the truthful outcome. If the truthful report yields the lexicographical optimal outcome, then clearly, there can be no responsively better outcome because a responsively better outcome is also a lexicographically better outcome. Now assume that the truthful outcome is lexicographically not the best. Then, there exists another allocation that is achievable that is lexicographically better. Consider the first (most preferred) item that is in the lexicographically more preferred outcome but not in the truthful outcome. We compute an allocation and partial preference of the manipulator that does not include any less preferred items. For the remaining items, we simply append them in the manipulator's preference list in decreasing order of preference. We claim that the outcome responsively dominates the truthful outcome from the manipulating agent's truthful preferences. The argument is technical and long and is omitted due to lack of space. \square

Example 6. *Let $n = 3, m = 7, \pi = 1231231$, and*

$\succ_1: o_1, o_2, o_3, o_4, o_5, o_6, o_7$

$\succ_2: o_2, o_1, o_5, o_3, o_4, o_6, o_7$

$\succ_3: o_2, o_3, o_4, o_1, o_5, o_7, o_6$

We have $O_1 = \{o_1, o_4, o_6\}$. At the first step, o_1 is added to O^ ; at the next step, o_2 is not added since $\{o_1, o_2\}$ is not feasible; at the next step, o_3 is added; as $O^* = \{o_1, o_3\}$ lexicographically dominates O_1 , we exit the loop, initialize \succ'_1 to*

(o_3, o_1) , and complete it into $\succ'_1: o_3, o_1, o_2, o_4, o_5, o_6, o_7$ which leads agent 1 to obtain $\{o_1, o_3, o_6\}$.

Conclusion

In this paper, we showed that computing a best response under sequential allocation to maximize additive utility is NP-hard. We also gave a contrasting results that manipulating sequential allocation under binary utilities and also with respect to responsive preferences is easy. Our NP-hardness result does not involve a constant number of agents. It remains an interesting open problem whether manipulating sequential allocation with respect to cardinal utilities is NP-hard when the number of agents is three or some other constant.

Acknowledgement This work is partly supported by the project ANR-14-CE24-0007-01 “CoCoRiCo-CoDec”.

References

- Aziz, H.; Gaspers, S.; Mackenzie, S.; Mattei, N.; Narodytska, N.; and Walsh, T. 2015a. Manipulating the probabilistic serial rule. In *Proc. of 14th AAMAS Conference*, 1451–1459.
- Aziz, H.; Gaspers, S.; Mackenzie, S.; and Walsh, T. 2015b. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71–92.
- Aziz, H.; Walsh, T.; and Xia, L. 2015. Possible and necessary allocations via sequential mechanisms. In *Proc. of 24th IJCAI*, 468–474.
- Berman, P.; Karpinski, M.; and Scott, A. D. 2003. Approximation hardness of short symmetric instances of MAX-3SAT. *Electronic Colloquium on Computational Complexity (ECCC)* (049).
- Bogomolnaia, A., and Moulin, H. 2004. Random matching under dichotomous preferences. *Econometrica* 72(1):257–279.
- Bouveret, S., and Lang, J. 2011. A general elicitation-free protocol for allocating indivisible goods. In *Proc. of 22nd IJCAI*, 73–78. AAAI Press.
- Bouveret, S., and Lang, J. 2014. Manipulating picking sequences. In *Proc. of 21st ECAI*, 141–146.
- Brams, S. J., and King, D. L. 2005. Efficient fair division: Help the worst off or avoid envy? *Rationality and Society* 17(4):387–421.
- Brams, S. J., and Straffin, P. D. 1979. Prisoners’ dilemma and professional sports drafts. *The American Mathematical Monthly* 86(2):80–88.
- Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Kalinowski, T.; Narodytska, N.; Walsh, T.; and Xia, L. 2013. Strategic behavior when allocating indivisible goods sequentially. In *Proc. of 27th AAAI Conference*, 452–458. AAAI Press.
- Kalinowski, T.; Narodytska, N.; and Walsh, T. 2013. A social welfare optimal sequential allocation procedure. In *Proc. of 22nd IJCAI*, 227–233. AAAI Press.

Kohler, D. A., and Chandrasekaran, R. 1971. A class of sequential games. *Operations Research* 19(2):270–277.

Levine, L., and Stange, K. E. 2012. How to make the most of a shared meal: Plan the last bite first. *The American Mathematical Monthly* 119(7):550–565.

Tominaga, Y.; Todo, T.; and Yokoo, M. 2016. Manipulations in two-agent sequential allocation with random sequences. In *Proc. of 15th AAMAS Conference*. IFAAMAS. 141–149.