ECAI 2014
T. Schaub et al. (Eds.)

doi:10.3233/978-1-61499-419-0-141

Manipulating picking sequences

Sylvain Bouveret[1] and **Jérôme Lang**[2]

Abstract. Picking sequences are a natural way of allocating indivisible items to agents in a decentralized manner: at each stage, a designated agent chooses an item among those that remain available. We address the computational issues of the manipulation of picking sequences by an agent or a coalition of agents. We show that a single agent can compute an optimal manipulation in polynomial time. Then we consider several notions of coalitional manipulation; for one of these notions, we show that computing an optimal manipulation is easy. We temper these results by giving a nontrivial upper bound on the impact of manipulation on the loss of social welfare.

1 INTRODUCTION

We study a very simple protocol for allocating indivisible goods to agents. The *picking sequence protocol* works as follows: we define a sequence of agents, and each agent is asked to take in turn one object among those that remain. For example, according to sequence $ABCCBA$, agent A will choose first, then agent B will pick one object, then C (two objects), and so on. This simple protocol is used in a lot of everyday life situations (allocating courses to students, initial resources in some board games...). Its simplest version, namely, the strict alternation protocol for two agents (*e.g.*, $ABABABAB$) has been studied first by Kohler and Chandrasekaran [9] in a game-theoretic setting, and then further by Brams and Taylor [4], who also pay attention to balanced alternation (*e.g.*, $ABBABAAB$) and Brams and King [3] for characterizing efficient allocations in a centralized setting. Budish and Cantillon [5] study a variant of the model (course allocation to students with a randomized version of a picking sequence) and show that not only it is manipulable in theory, but that it also is manipulated by students in practice. Picking sequences were formally studied in a more general and systematic way by Bouveret and Lang [2], and further by Kalinowski *et al.* [8] who give a game-theoretic study of picking sequences, and Kalinowski *et al.* [7] who (among other results) prove that for a plausible set of criteria, strict alternation is the best picking sequence for two agents.

In this paper, we study this protocol from the point of view of single-agent and coalitional manipulation. The strategical issues of picking sequences have already been studied by Kohler and Chandraesekaran [9], who prove that the subgame perfect Nash equilibrium can be computed in by reversing the policy and preference orderings. Kalinowski *et al.* [8] extend this result to any two-agent picking sequence, and investigate the computational complexity of computing a subgame perfect Nash equilibrium for more than two agents. These papers give a game-theoretic study of picking sequences: more precisely, they focus on the characterization and the computation of subgame perfect Nash equilibria. In this paper, we use a different approach. We view manipulation in picking sequences exactly as manipulation in voting. Voting theory, and especially computational social choice, has devoted a lot of attention to the manipulation of voting rules by a single deviating agent, or by a coalition of deviating agents (see [6] for a recent survey); the assumption on both cases is that the votes of the non-manipulators are known. This approach to manipulation in picking sequences remains largely unexplored, if we except some previous results [2].

The paper is organized as follows. We introduce some formal background in Section 2. Then we study manipulation by a single agent (Section 3) and a coalition of agents (Section 4) under various assumptions. In Section 5 we briefly study the price of manipulation in a two-agent setting, that is, the worst-case loss of social welfare caused by one agent acting strategically. We conclude in Section 6.

2 BACKGROUND AND NOTATIONS

$\mathcal{N} = \{A, B, \dots\}$ is a set of n agents and $\mathcal{O} = \{1, \dots, m\}$ a set of m objects. Each agent i is equipped with a (private) preference relation $\succeq_i$, which is a weak order (transitive and complete relation) on $2^{\mathcal{O}}$. The restriction of $\succeq_i$ to $\mathcal{O}$ is denoted by $\trianglerighteq_i$. We write $\succ$ (resp. $\triangleright$) to denote the strict part of $\succeq$ (resp. $\trianglerighteq$). $\succeq_i$ is *(weakly) separable*, if for any $S, S' \subseteq \mathcal{O}$ and $o, o' \in \mathcal{O} \setminus (S \cup S')$, $S \cup \{o\} \succeq_i S \cup \{o'\} \Leftrightarrow S' \cup \{o\} \succeq_i S' \cup \{o'\}$.

A *policy* π is a sequence of m agents. For any agent i, we write $ps(\pi, i)_1, \dots, ps(\pi, i)_{r(i)}$ to denote the $r(i)$ successive *picking stages* of agent i. For all k, we will also denote by $PS(\pi, i)_k$ the number of picking stages of agent i until stage k. A *(deterministic, simple) picking strategy* for agent i is a function $\sigma_i : 2^{\mathcal{O}} \to \mathcal{O}$, specifying which object $\sigma_i(O)$ agent i should take when the set of remaining objects is O (see Section 3 for more about strategies).

Given a set of agents C, a *joint strategy* for C is a function σ_C mapping each agent $i \in C$ to a given strategy σ_i. Given a joint strategy σ concerning all the agents, we will denote by σ_{-i} the joint strategy of all the agents but i, and similarly σ_{-C} the joint strategy of all the agents but the ones in C. Given two strategies σ and τ concerning different agents, $\sigma \cdot \tau$ will denote the joint strategy built out from σ and τ. Finally, given a joint strategy σ concerning all the agents and a policy π, we will denote by $\mathcal{O}_i(\pi, \sigma)$ the set of objects obtained by agent i in picking sequence π if every agent j follows strategy σ_j.

We make use of the following notation for allocations resulting form a picking sequence: $[O_1| \dots |O_n]$ is the allocation where agent i has set of objects O_i. Moreover, we often omit curly brackets for sets. For instance, if $n = 3$ and $m = 7$, $[1234|5|67]$ is the allocation giving $\{1, 2, 3, 4\}$ to agent A, $\{5\}$ to agent B and $\{6, 7\}$ to agent C.

3 MANIPULATION BY A SINGLE AGENT

Clearly, the only strategyproof picking sequences are those where each agent acts in a single "picking row", without alternation, that

[1] LIG – Grenoble INP; sylvain.bouveret@imag.fr
[2] LAMSADE – Univ. Paris-Dauphine; lang@lamsade.dauphine.fr

is, sequences of the form $\sigma = a_{i_1}^{m_1} a_{i_2}^{m_2} \ldots a_{i_k}^{m_k}$, where $a_{i_1}, \ldots, a_{i_k}$ are all different agents, and $m_1 + \ldots + m_k = m$.[3]

Bouveret and Lang ([2], Proposition 7) show that finding a manipulation for an n-agent picking sequence can be polynomially reduced to finding a manipulation for a 2-agent picking sequence. Therefore, without loss of generality, we consider in this Section that we have only two agents $\{A, B\}$, where A will be the manipulating agent.

A standard approach to manipulation in voting is to consider that the manipulating agent has a complete knowledge about the other agents' votes. We make a similar assumption here: *A has a complete knowledge of B's picking strategy.* We now make a stronger assumption about B's strategy: we assume that *B picks, at each stage, the best object among the remaining ones, according to a (real or virtual) linear order $\rhd'_B$, that A knows.* Such picking strategies are said to be *deterministic* and *simple*.

Assuming that B's strategy is deterministic, simple, and known to A, implies a loss of generality. First, if B has nonseparable preferences, he may pick items in a perfectly rational way according to a complex choice function not rationalizable by a weak order over objects: for instance, if B is interested in getting 1 and 2 together, but not interested in getting only one of them, then he could pick 1 if he has two more picking stages and the remaining objects are $\{1, 2, 3, 4\}$ but 3 if the remaining objects are $\{1, 3, 4\}$. Thus, in general, a deterministic picking strategy for B would be an arbitrary function from $2^{\mathcal{O}}$ to $\mathcal{O}$.[4] Second, even if B has separable preferences, B can be indifferent between two (or more) given objects, and in this case, it is not clear (even to him) which of the two it will pick first. (A similar phenomenon occurs in voting: nonmanipulators may be indifferent between some candidates, yet the manipulator is assumed to know how they will rank them.) It could even make sense to assume that B acts according to a mixed strategy.

While it would make sense to study manipulation with such mixed and/or complex strategies of B, we leave them for further research and assume here that

(Hyp) A knows B's picking strategy σ_B, and this strategy is a deterministic, simple picking strategy.

where a deterministic, simple strategy σ_B is represented by a ranking $\rhd'_B$ over $\mathcal{O}$, such that B always picks the preferred object, with respect to $\rhd'_B$, among those that remain available. From now on we will simply say "picking strategy" for "deterministic, simple picking strategy". It must be noted that σ_B is not necessarily sincere, and it does not need to be: $\rhd'_B$ can be but is not necessarily related to agent B's true preference relation $\succeq_B$. However, to avoid overloading notation, and since we will not deal at all with agent B's true preferences, $\rhd'_B$ will simply be denoted by $\rhd_B$. In the following, we will denote by $\sigma^{\rhd}$ the sincere picking strategy defined (unambiguously) from $\rhd$.

Because A has a complete knowledge of B's strategy, it is enough for her to choose a simple deterministic strategy as well, which amounts to choosing the best possible set of objects she can get.

Now we discuss various assumptions about A's preferences. In the simplest case, we assume that A's preferences are additive with no ties on single objects. This means that A's preferences can be represented succinctly by a utility function over single goods $u_A : G \to \mathbb{R}^+$. However, we shall see in Subsection 3.1 that A's optimal strategy does not depend on the values of u *but only on the corresponding ranking $\rhd_A$ on objects* – just as what we need to know for the sincere strategy. In Subsection 3.2 we assume that A's preferences are additive with possible ties on single objects. Lastly, in Section 3.3 we will consider manipulators with possibly non-separable preferences.

[3] Such non-alternating sequences are in fact a kind of *sequential dictatorships*. In settings where agents get only one object, sequential dictatorships are the only strategyproof resource allocation mechanisms satisfying a set of mild properties [10].

[4] Even modelling a deterministic strategy for B as a function from $2^{\mathcal{O}}$ to $\mathcal{O}$ could be a loss of generality: in some contexts, B's strategy could also depend in the *order* in which A has picked her objects.

3.1 The manipulator has additive preferences without indifferences between single objects

From what we said above (Hyp + A has additive preferences without indifferences), a two-agent picking sequence manipulation problem for manipulator A is a quadruple $\langle \mathcal{O}, u_A, \rhd_B, \pi \rangle$ where:

- $\mathcal{O}$ is a set of m objects;
- $u_A : \mathcal{O} \to \mathbb{R}^+$ is A's utility function over single objects, and verifies $u(o) \neq u(o')$ if $o \neq o'$.
- $\rhd_B$ is a ranking over $\mathcal{O}$ (succinct representation of B's simple deterministic picking strategy);
- $\pi \in \{A, B\}^m$ is a picking sequence.

A's utility over subsets of goods (denoted by u_A as well) is defined by $u_A(S) = \sum_{o \in S} u_A(o)$, for any $S \subseteq \mathcal{O}$.

Let $\rhd_A$ be the ranking over $\mathcal{O}$ induced by u_A, that is, $o \rhd_A o'$ if and only if $u_A(o) > u_A(o')$. If $\rhd_A$ is induced by u_A we also say that u_A is *compatible* with $\rhd_A$.

Let $P = \langle \rhd_A, \rhd_B \rangle$. Without loss of generality we assume that $1 \rhd_A 2 \rhd_A \ldots \rhd_A m$. σ^P will be the joint sincere strategy $\sigma^{\rhd_A} \cdot \sigma^{\rhd_B}$, and thus $\langle \mathcal{O}_A(\sigma^P, \pi), \mathcal{O}_B(\sigma^P, \pi) \rangle$will denote the allocation resulting from the sincere picking of agents A and B according to P and π.

We now define manipulation. We say that a subset of objects O is *achievable* for A if there is a strategy σ such that $O \subseteq \mathcal{O}_A(\sigma \cdot \sigma^{\rhd_B}, \pi)$. In other words, A can obtain all the objects from O by playing according to σ. A manipulation strategy σ is *successful* if $\mathcal{O}_A(\sigma \cdot \sigma^{\rhd_B}, \pi) \rhd_A \mathcal{O}_A(\sigma^P, \pi)$ (in other words, A obtains a better set by playing according to σ than if it had played sincerely).

We already know from Proposition 7 in [2] that it can be checked in polynomial time whether a given set O is achievable. An important problem is to determine whether there is a successful strategy. It turns out that not only can we solve this problem in polynomial time, but we are also able to find the strategy giving the best achievable subset in polynomial time as well, using Algorithm 1.

Algorithm 1: Best achievable subset.

```
input  : A policy π, a ranking ▷_B.
output: The best achievable subset for agent A.
O ← ∅;
j ← 1;
for k ← 1 to m do
    find the smallest i ≥ j such that O ∪ {i} is achievable;
    O := O ∪ {i};
    j := i;
return O;
```

Because it can be checked in polynomial time whether a given set is achievable, Algorithm 1 works in polynomial time. Before proving that it indeed returns an optimal strategy, we give some examples.

Example 1 $m = 12$; $\pi = ABABABABABABAB$;

$\rhd_A$:	1	2	3	4	5	6	7	8	9	10	11	12
$\rhd_B$:	3	2	6	5	4	10	8	11	1	9	7	12

The sincere picking strategy for A (and B) leads to A getting 124789. *Let us apply the algorithm.*

- 1 *is achievable;* $S = 1$*;*
- 12 *is achievable;* $S = 1, 2$*;*
- 123 *is not achievable;* $S = 1, 2$*;*
- 124 *is achievable;* $S = 1, 2, 4$*;*
- 1245 *is achievable;* $S = 1, 2, 4, 5$*;*
- 12456 *is not achievable;* $S = 1, 2, 4, 5$*;*
- 124567 *is achievable;* $S = 1, 2, 4, 5, 7$*;*
- 1245678 *is achievable;* $S = 1, 2, 4, 5, 7, 8$*; stop and return* S.

We will soon prove that this is indeed the best achievable set of objects for A, irrespective of the choice of the utility function satisfying the requirement of this Subsection (namely, that A has an additive utility function with all weights on single objects being different).

For any s, we denote B's sth preferred object by $\mu(s)$: μ is the permutation of $\{1, \ldots, m\}$ such that $\mu(1) \rhd_B \ldots \rhd_B \mu(m)$. Moreover, let $B(s) = \{\mu(1), \ldots, \mu(s)\}$ be the set of B's s most preferred items. Finally, for any $X \subseteq \mathcal{O}$, let $Cl(s, X) = B(s) \cap X$. We know from Proposition 8 in [2] that there exists a successful picking strategy σ for $X \subseteq \mathcal{O}$ if and only if for every picking stage s, $PS(\pi, A)_s \geq |Cl(s, X)|$. Also, we know (again Proposition 7 in [2]) that if O is achievable, then the *standard picking strategy*, in which A picks items in O according to their increasing ranking in $\rhd_B$, is successful. Such a strategy will be denoted by $\sigma_{st(O)}$.

Example 2 (continued)

- $\mu(1) = 3; \mu(2) = 2$*; etc.*
- $B(1) = \{3\}$*;* $B(2) = \{3, 2\}$*; etc.*
- *Let us check that* 1245 *is achievable:*
  - $Cl(1, 1245) = \emptyset$*;* $PS(\pi, A)_1 = 1$*;*
  - $Cl(2, 1245) = \{2\}$*;* $PS(\pi, A)_2 = 1$*;*
  - $Cl(3, 1245) = \{2\}$*;* $PS(\pi, A)_3 = 2$*;*
  - $Cl(4, 1245) = \{2, 5\}$*;* $PS(\pi, A)_4 = 2$*;*
  - $Cl(5, 1245) = \{2, 5, 4\}$*;* $PS(\pi, A)_5 = 3$*;*
  - $Cl(6, 1245) = \{2, 5, 4\}$*;* $PS(\pi, A)_6 = 3$*;*
  - $Cl(7, 1245) = \{2, 5, 4\}$*;* $PS(\pi, A)_7 = 4$*;*
  - $Cl(8, 1245) = \{2, 5, 4\}$*;* $PS(\pi, A)_8 = 4$*;*
  - $Cl(9, 1245) = \{1, 2, 5, 4\}$*;* $PS(\pi, A)_9 = 5$*; etc.*
- *Let us check that* 123 *is not achievable:*
  - $Cl(2, 123) = \{2, 3\}$*;* $PS(\pi, A)_2 = 1$.
- *the standard picking strategy for* 1245 *is* $\sigma(1) = 2; \sigma(2) = 5; \sigma(3) = 4; \sigma(1) = 1$*, which we abbreviate in* $\langle 2, 5, 4, 1\rangle$*. For* 124578 *it is* $\langle 2, 5, 4, 8, 1, 7\rangle$.

Lemma 1 *Assume that O and $O' \neq O$ are achievable. Let $i = \min((O' \setminus O) \cup (O \setminus O'))$ and assume that $i \in O'$. Let j be B's most preferred item in $O \setminus O'$. Then $O[i \leftrightarrow j] = (O \cup \{i\}) \setminus \{j\}$ is achievable.*

Proof: In the following proof, we will refer to Proposition 7 in [2] as (P). We consider two cases, according to B's preference between i and j:

Case 1: $j \rhd_B i$**.** Because $j \rhd_B i$, we have that, for every picking stage s, $Cl(s, O[i \leftrightarrow j]) \subseteq Cl(s, O)$. Therefore, by (P), $O[i \leftrightarrow j]$ is achievable.[5]

Case 2: $i \rhd_B j$**.** Assume $O[i \leftrightarrow j]$ is not achievable. Then, by (P), there is a picking stage s such that

$$PS(\pi, A)_s < |Cl(s, O[i \leftrightarrow j])|. \qquad (1)$$

Let s^* be such that $\mu(s^*) = j$. We consider two cases.

Case 2.a: $s \leq s^*$**.** Since j is B's most preferred item in $O \setminus O'$, it holds that every item $l \in O$ such that $l \rhd_B j$ also belongs to O'. Obviously, that also holds for every item $l \in O[i \leftrightarrow j]$ such that $l \rhd_B j$. Hence $B(t) \cap O[i \leftrightarrow j] \subseteq B(t) \cap O'$ for all $t < s^*$. This equation can also be extended to $t = s^*$ by using the fact that $\mu(s^*) = j$ and that j neither belongs to O' nor to $O[i \leftrightarrow j]$. This in turn can be rewritten as $Cl(t, O[i \leftrightarrow j]) \subseteq Cl(t, O')$ for all $t \leq s^*$. Using this equation for $t = s$ together with Equation (1) leads to $PS(\pi, A)_s < |Cl(s, O')|$, which proves, using (P), that O' is not achievable. Contradiction.

Case 2.b: $s > s^*$**.** Since $\mu(s^*) = j$, j belongs to $B(s^*)$ and hence to $B(s)$. Since $i \rhd_B j$, i also belongs to $B(s^*)$ and hence to $B(s)$. Therefore, $B(s) \cap O = B(s) \cap O[i \leftrightarrow j]$, which, once again, can be rewritten as $Cl(s, O) = Cl(s, O[i \leftrightarrow j])$. Hence, by Equation (1), it holds that $PS(\pi, A)_s < |Cl(s, O)|$, which proves, using (P), that O is not achievable. Contradiction. ■

Example 3 (continued) $O = 124789$*;* $O' = 1245$*;* $i = 5$*;* $O \setminus O' = 789$*;* $j = 8$. $O[5 \leftrightarrow 8] = 124579$. *Lemma 1 says that* 124579 *is achievable. We are in the case* $i \rhd_B j$.

Proposition 1 *Algorithm 1 returns the best achievable subset for A.*

Proof: Assume not. Let O' be an optimal achievable subset, and O the subset returned by the algorithm. Let $i = \min((O' \setminus O) \cup (O \setminus O'))$. By construction of O, we must have $i \in O$. Now, by Lemma 1, there exists $j \in O'$, $j > i$, such that $(O' \cup \{i\}) \setminus \{j\}$ is achievable. Now, $i \rhd_A j$, *i.e.*, $u_A(i) > u_A(j)$, therefore, $u_A(O' \cup \{i\}) \setminus \{j\}) > u_A(O')$, which contradicts the optimality of O'. ■

Corollary 1 *An optimal manipulation for $(\mathcal{O}, u_A, \rhd_B, \pi)$ can be computed in polynomial time.*

Another consequence of Proposition 1 is the uniqueness of the best achievable subset for A. Thus, even if there may be several optimal manipulations, they are equivalent in the sense that the outcome for A is the same for all.

Importantly, note that the proof of Proposition 1 does not depend on the values of u_A (provided, as assumed at the beginning of the Subsection, that $o \neq o'$ implies $u_A(o) \neq u_A(o')$) but only on the order $\rhd_A$. We state this as a formal result:

Proposition 2 *The optimal manipulations for A are the same for any utility function u_A compatible with $\rhd_A$.*

3.2 The manipulator has additive preferences with possible indifferences between single objects

Now, a two-agent picking sequence manipulation problem for manipulator A is a quadruple $\langle \mathcal{O}, u_A, \rhd_B, \pi\rangle$ where:

[5] Even if we don't need it for the proof, the picking strategy obtained from the standard picking strategy $\sigma_{st(O)}$ by replacing j by i is successful – note that it does not necessarily correspond to the standard picking strategy for $O[i \leftrightarrow j]$.

- $\mathcal{O}$ is a set of m objects;
- $u_A : \mathcal{O} \to \mathbb{R}^+$ is A's utility function over single objects.
- $\rhd_B$ is a ranking over $\mathcal{O}$ (succinct representation of B's simple deterministic picking strategy);
- $\pi \in \{A, B\}^m$ is a picking sequence.

Now, the preference relation over single objects induced from u_A is a weak order $\unrhd_A$ over $\mathcal{O}$, defined by $o \unrhd_A o'$ if and only if $u_A(o) \geq u_A(o')$. Let $\sim_A$ (respectively, $\rhd_A$) be the indifference (resp. strict preference) relation associated with $\unrhd_A$, defined by $o \sim_A o'$ if and only if $o \unrhd_A o'$ and $o' \unrhd_A o$ (respectively, $o \unrhd_A o'$ and $o' \not\unrhd_A o$).

Now let $\rhd'_A$ be the linear order on $\mathcal{O}$ refining $\unrhd_A$ and defined by: $o \rhd'_A o'$ if and only if $o \rhd_A o'$ or ($o \sim_A o'$ and $o \rhd_B o'$). For example, if $1 \rhd_A 2 \sim_A 3 \sim_A 4 \rhd_A 5$ and $3 \rhd_B 4 \rhd_B 1 \rhd_B \rhd_B 5 \rhd_B 2$, then $1 \rhd'_A 3 \rhd'_A 4 \rhd'_A 2 \rhd'_A 5$. Let u'_A be a utility function on $\mathcal{O}$ compatible with $\rhd'_A$. We claim that an optimal achievable set of objects for A can be computed as follows.

Proposition 3 *The (unique) optimal achievable subset for* $(\mathcal{O}, u'_A, \rhd_B, \pi)$ *is a (non necessarily unique) optimal achievable subset for* $(\mathcal{O}, u_A, \rhd_B, \pi)$.

Proof: Let Y be the optimal achievable set of objects returned by the resolution of the manipulation problem for $(\mathcal{O}, u'_A, \rhd_B, \pi)$. Assume that Y is not optimal for $(\mathcal{O}, u_A, \rhd_B, \pi)$, that is, there is an achievable set Z such that $u_A(Z) > u_A(Y)$. Let δ be such that $0 < \delta \leq |u_A(O) - u_A(O')|$ for all subsets $O, O' \subseteq \mathcal{O}$ such that $u(O) \neq u(O')$, and let $\varepsilon < \frac{\delta}{m}$. Now, let u''_A be the following utility function: $\forall i \in \mathcal{O}, u''_A(i) = u_A(i) + \varepsilon q(i)$, where $q(i) = |\{j \mid i \sim_A j \text{ and } i \rhd_B j\}|$. The following facts hold: (i) u''_A is compatible with $\rhd'_A$, and (ii) $u''_A(Z) > u''_A(Y)$. To prove (i), let i and j be two objects. We consider two cases. (a) $i \sim_A j$: then $u''_A(i) > u''_A(j)$ iff $q(i) > q(j)$ iff $i \rhd_B j$ iff $i \rhd'_A j$. (b) $i \not\sim_A j$ (assume wlog $i \rhd_A j$): then $u''_A(i) - u''_A(j) = u_A(i) - u_A(j) + (q(i) - q(j))\varepsilon > u_A(i) - u_A(j) - m\varepsilon \geq u_A(i) - u_A(j) - \delta > 0$. (ii) can be proved as follows: $u''_A(Z) - u''_A(Y) \geq u_A(Z) - u_A(Y) - m\varepsilon > u_A(Z) - u_A(Y) - \delta > 0$. (i) and (ii) together prove that Y cannot be the optimal achievable set of objects for $(\mathcal{O}, u''_A, \rhd_B, \pi)$, and also for $(\mathcal{O}, u'_A, \rhd_B, \pi)$, since u'_A and u''_A are both compatible with $\unrhd_A$: contradiction. ■

From Corollary 1 and Proposition 3 we get:

Corollary 2 *An optimal manipulation for* $(\mathcal{O}, \rhd'_A, \rhd_B, \pi)$ *can be computed in polynomial time.*

Also, we have a result analogous to Proposition 2: the optimal achievable subset, and the picking strategy, is optimal irrespective of the choice of the utility function u_A extending $\unrhd_A$.

3.3 The manipulator has non-additive preferences

Assume now that the manipulator A no longer has additive preferences. One of the simplest forms of non-additive preferences are (unrestricted) dichotomous monotonic preferences: there is a set of objects $Good_A \subseteq \mathcal{O}$ such that (a) $Good_A$ is upward closed, that is, if $S \subseteq S'$ and $S \in Good_A$ then $S' \in Good_A$, and (b) A equally likes all subsets in A and equally dislikes all subsets in $2^{\mathcal{O}} \setminus Good_A$, that is, $S \unrhd_A S'$ if and only if ($S' \in Good_A$ implies $S \in Good_A$). We know (see for instance [1]) that a dichotomous monotonic preference relation can be represented succinctly by a positive (negation-free) propositional formula φ_A of the language $\mathcal{L}_{\mathcal{O}}$ constructed from a set of propositional symbols isomorphic to $\mathcal{O}$. For instance, $o_1 \vee (o_2 \wedge o_3)$ means that any set containing o_1 or both o_2 and o_3 is good for A: $\{o_1, o_2, o_3\} \sim_A \{o_1, o_2\} \sim_A \{o_1, o_3\} \sim_A \{o_1\} \sim_A \{o_2, o_3\} \rhd_A \sim_A \{o_2\} \sim_A \{o_3\} \sim_A \emptyset$.

Thus, a two-agent picking sequence manipulation problem for manipulator A with dichotomous monotonic preferences is a triple $\langle \mathcal{O}, \varphi_A, \rhd_B, \pi \rangle$ where $\mathcal{O}$, $\rhd_B$ and π are as usual, and φ_A is a positive propositional formula of $\mathcal{L}_{\mathcal{O}}$.

We say that a picking strategy for A is *successful* if it gives her a set of objects in $Good_A$. Since all sets in $Good_A$ are equally good, optimal picking strategies coincide with successful strategies provided that there exists at least one (and with all strategies otherwise).

Proposition 4 *Deciding whether a manipulation problem* $\langle \mathcal{O}, \varphi_A, \rhd_B, \pi \rangle$ *has a successful picking strategy is* NP*-complete, even if* π *is the alternating sequence.*

Proof: Membership is obvious (guess the picking strategy and apply it). Hardness follows by reduction from SAT. Let $\alpha = c_1 \wedge \ldots \wedge c_k$ be a propositional formula under conjunctive normal form over a set of propositional symbols $\{x_1, \ldots, x_p\}$. Define the following instance of a manipulation problem $\langle \mathcal{O}, \varphi_A, \rhd_B, \pi \rangle$:

- $\mathcal{O} = \{o_1, o'_1, \ldots, o_p, o'_p\}$;
- for every clause c_i of α, let c'_i be the clause obtained by replacing every positive literal x_i by o_i and negative literal $\neg x_i$ by o'_i; let α' be the conjunction of all clauses c'_i and finally, let $\varphi_A = \alpha' \wedge \bigwedge_{i=1}^{p} (o_i \vee o'_i)$;
- $\pi = (AB)^p$
- $\rhd_B = o_1 \rhd o'_1 \rhd o_2 \rhd o'_2 \rhd \ldots \rhd o_p \rhd o'_p$.

If φ is satisfiable then let $\omega \models \varphi$; consider the picking strategy in which, at her ith picking stage A picks o_i if ω assigns x_i to true and o'_i if ω assigns x_i to false (and then B will pick o'_i if A has picked o_i, and o_i if A has picked o'_i). The resulting set of objects will be exactly $S = \{o_i | \omega \models x_i\} \cup \{o'_i | \omega \models \neg x_i\}$, and since $\omega \models \alpha$, we have that S satisfies α'; moreover, clearly S satisfies $o_i \vee o'_i$ for each i, therefore, S satisfies φ_A.

Conversely, assume that A has a picking strategy that leads to a set of objects S satisfying φ_A. Because S contains one of o_i and o'_i for each i, and because $|S| = p$, S contains *exactly* one of o_i and o'_i for each i. Let ω be the interpretation over $\{x_1, \ldots, x_p\}$ defined by $\omega \models x_i$ if $o_i \in S$ and $\omega \models \neg x_i$ if $o_i \notin S$. Because S satisfies α', we have that $\omega \models \alpha$, that is, α is satisfiable. ■

As a consequence, more generally, deciding whether a manipulation problem (with arbitrary, compactly represented preferences) has a successful picking strategy is NP-hard.

4 COALITIONAL MANIPULATION

Voting theory not only focuses on single-agent manipulation but also on *joint (or coalitional) manipulation*, where a group of voters collude to get a better outcome for themselves. It is assumed that they can fully communicate and have full knowledge of the others' votes.

However, there is a significant difference with voting: the outcome of a vote is the same for all agents, whereas in fair division agents get different shares and are thus allowed to make trades after the allocation has been made. We thus consider three different notions of manipulation. The first two do not need any particular assumption about voters' preferences. The first one says that a manipulation is a combination of picking strategies whose outcome Pareto-dominates (for the manipulating coalition) the outcome of the sincere picking strategy; it does not allow any posterior trading nor compensation. The second one is also based on Pareto-dominance but allows agents to trade items after the allocation has been done. The third one assumes that the manipulators' preferences are represented by transferable utilites, and allows both trading *and* monetary transfers after the allocation has been done. Before giving the formal definition we give a few examples. *In all cases, we have three agents A, B, C, and the manipulation coalition consists of A and B.*

Example 4 $\pi = ABCABC$. *No post-allocation trade is allowed.*

$$\rhd_A: 125436;\ \ \rhd_B: 135246;\ \ \rhd_C: 234156$$

Sincere picking leads to [15|34|26]. *A and B manipulating alone cannot do better: their best responses to the other two players' sincere strategies is their sincere strategy. However, if they cooperate, then they both can do better: A start by picking 2, then B picks 3, C picks 4, A picks 1, B picks 5 and finally C picks 6. The final allocation is* [12|35|46], *which (strongly) Pareto-dominates* [15|34|26]. *Note that it is crucial that A and B communicate beforehand and trust each other, for after A has picked 2, B can betray A and pick 1 instead of 5, resulting in the final allocation* [25|14|36], *which may be better for B then the joint strategy agreed upon if she values $\{1,4\}$ more than $\{3,5\}$, but for A is worse than the sincere allocation.*

Example 5 $\pi = ABCABCABC$. *Post-allocation exchange of goods is allowed. Monetary transfers are not allowed.*

$$\rhd_A: 123456789;\ \ \rhd_B: 893456712;\ \ \rhd_C: 123897456$$

Sincere picking leads to [134|589|267]. *A and B manipulating alone cannot do better. They also cannot do better if they are not allowed to exchange goods (we will see later how to check this). However, if they cooperate and are allowed to exchange goods, then A can start by picking 1, then B picks 2, C picks 3, A picks 8, B picks 9, C picks 7, A picks 4, B picks 5 and C picks 6, leading to* [148|259|367]. *then A and B exchange 2 and 8, leading to* [124|589|367], *which Pareto-dominates* [134|589|367] *for $\{A, B\}$.*

Example 6 $\pi = ABCABCABC$. *Post-allocation exchange of goods is allowed. Monetary transfers are not allowed.*

$$\rhd_A: 123456789;\ \ \rhd_B: 345916782;\ \ \rhd_C: 123897459$$

Assume that B prefers 459 to 358. Sincere picking leads to [147|358|269]. *If A and B cooperate they can get* [147|259|368], *then swap 2 and 4, leading to* [127|459|368]*: both agents are better off. This, of course, depends on some extra information, that is, the manipulators' preferences over the full power set.*

Example 7 $\pi = ABCABCABC$. *Post-allocation exchange of goods is allowed. Monetary transfers are allowed.*

$$\rhd_A: 123456789;\ \ \rhd_B: 987654321;\ \ \rhd_C: 123897459$$

Assume that A and B have additive preferences, that correspond to the amount of money they are willing to pay to get the items, and that

- $u_A(1) = 14$; $u_A(2) = 13$; $u_A(3) = 12$; $u_A(4) = 11$; $u_A(5) = 10$; $u_A(6), u_A(7)... \leq 5$;
- $u_B(9) = 10$; $u_B(8) = 9$; $u_B(7) = 8$; $u_B(6) = 7$; $u_B(5) = 6$; *the rest does not matter.*

Sincere picking leads to [125|789|346]. *If A and B cooperate they can get* [147|259|368], *then B gives 2 and 5 to A, A gives 7 to B together with some amount of money. Both are strictly better off. This needs transferable utilities*

These examples illustrate three different ways of defining what makes a coalition better off. More formally:

Definition 1 *Let $\mathcal{N}$ be a set of agents, π be a sequence, and $C \subset \mathcal{N}$ be a coalition of agents. Moreover, let σ_C and σ'_C be two joint strategies for C. We will say that :*

- σ_C Pareto-dominates σ'_C *(written $\sigma_C > \sigma'_C$) if:*
  - $\forall i \in C$, $\mathcal{O}_i(\pi, \sigma_C \cdot \sigma^*_{-C}) \succeq_i \mathcal{O}_i(\pi, \sigma'_C \cdot \sigma^*_{-C})$*;*
  - *this inequality is strict for at least one i.*
- σ_C Pareto-dominates with transfers σ'_C *(written $\sigma_C >_T \sigma'_C$) if there is a function $F : \bigcup_{i \in C} \mathcal{O}_i(\pi, \sigma_C \cdot \sigma^*_{-C}) \to C$ such that:*
  - $\forall i \in C$, $\{k \in \mathcal{O} \mid F(k) = i\} \succeq_i \mathcal{O}_i(\pi, \sigma'_C \cdot \sigma^*_{-C})$*;*
  - *this inequality is strict for at least one i.*

Finally, if we assume that each agent i (at least those from C) is equipped with a valuation function $v_i : 2^{\mathcal{O}} \to \mathbb{R}$, compatible with $\rhd_i$, we will say that:

- σ_C Pareto-dominates with transfers and side-payments σ'_C *(written $\sigma_C >_{TP} \sigma'_C$) if there is a function $F : \bigcup_{i \in C} \mathcal{O}_i(\pi, \sigma_C \cdot \sigma^*_{-C}) \to C$, and a function $p : C \to \mathbb{R}$ such that:*
  - $\sum_{i \in C} p_i = 0$*;*
  - $\forall i \in C$, $v_i(\{k \in \mathcal{O} \mid F(k) = i\}) + p(i) \geq v_i(\mathcal{O}_i(\pi, \sigma'_C \cdot \sigma^*_{-C}))$*;*
  - *this inequality is strict for at least one i.*

These definitions lead to three notions of successful strategies:

- σ_C is a *successful strategy* if $\sigma_C > \sigma^*_C$;
- σ_C is a *successful strategy with transfers* if $\sigma_C >_T \sigma^*_C$;
- σ_C is a *successful strategy with transfers and side-payments* if $\sigma_C >_{TP} \sigma^*_C$.

In the following, we will focus on the following problem:

CM-SIMPLE	
Given:	A set of agents $\mathcal{N}$, a sequence π, a coalition $C \subset \mathcal{N}$ with their preference relations $\succeq_i$ and a joint strategy σ_C
Question:	Is there a strategy σ'_C such that $\sigma'_C > \sigma_C$?

The variant with transfers ($\sigma'_C >_T \sigma_C$) and transfers with side-payments ($\sigma'_C >_{TP} \sigma_C$ – in this case, we need to add the coalition members' valuation functions v_i to the problem input) will be called respectively CM-TRANSFERS and CM-TRANSFERSWITHPAYMENTS. Note that if we want to know whether a successful strategy exists for a given setting, we just need to solve the latter problem with σ_C being the sincere strategy σ^*_C.

We start by considering manipulators with additive preferences.

Proposition 5 *An optimal manipulation for a coalition of agents M with transfers and side payments can be found in polynomial time.*

Proof: The possibility of side payments and exchanges imply that (a) in the optimal final allocation (after the exchanges), each object will be assigned to the agent who gives it the highest utility (or one of the agents who gives it the highest utility, in case there are several), and (b) the optimal joint picking strategy is the one that maximizes the utilitarian social welfare of the group of manipulators $\sum_{i \in M} v_i(S_i)$. (a) and (b) together imply that the optimal set S of objects for the group maximizes $\sum_{o \in S} \max_{i \in M} v_i(o)$. This is equivalent to solving a manipulation problem for a single manipulator with a weak order over objects $o \trianglerighteq o'$ iff $\max_{i \in M} v_i(o) \geq \max_{i \in M} v_i(o')$. Proposition 3 then guarantees that such an optimal manipulation can be found in polynomial time. ■

For coalitional manipulation without monetary transfers we have the following results; due to space limitations, proofs are omitted.[6]

Proposition 6 *Deciding if there exists a successful manipulation without transfers nor side payments is* NP*-complete,* even for two manipulators with additive preferences and no non-manipulator.

Proposition 7 *Deciding if there exists a successful manipulation with transfers and without side payments is* NP*-complete, even for two manipulators with additive preferences.*

Finally, in the case of non-additive preferences, Proposition 4 directly entails that CM-SIMPLE is NP-hard, for any set of non-additive preference relations $\succ_i$ represented in a compact way.

5 PRICE OF MANIPULATION

The results of Sections 3 and 4 can be seen as an argument against using picking sequences. However, we continue thinking that, in spite of this, picking sequences is one of the best protocol for allocating objects without prior elicitation, because of its simplicity. Moreover, we now temper the results about the easiness of manipulation by showing that, at least in some simple cases, the worst-case price of manipulation (that is, the loss of social welfare caused by one agent manipulating) is not significantly high. Note that, to define the price of manipulation properly, we need to deal with numerical preferences. A classical technique to translate ordinal preferences into utility functions is to use *scoring functions*, as in voting. Formally, a scoring function g is a non-decreasing function from $\{1, \dots, m\}$ to $\mathbb{R}$. $g(j)$ is the utility an agent i receives for an object ranked at position j $\rhd_i$. For each agent i, u_i is computed by summing the utilities $g(j)$ for each object i receives, using the same scoring function g.

Definition 2 *Let $P = \langle \rhd_A, \rhd_B, \dots \rangle$ be a preference profile, π be a sequence, and g be a scoring function. Let σ_A be a successful manipulating strategy for agent A. The* price of manipulation *for σ_A given (P, π, g) is the ratio:*

$$PM_{P,\pi,g}(\sigma_A) = \frac{\sum_{i \in N}(u_i(\mathcal{O}_i(\pi, \sigma_A \cdot \sigma^*_{-A})))}{\sum_{i \in N}(u_i(\mathcal{O}_i(\pi, \sigma^*_{\mathcal{N}})))}.$$

In other words, the price of manipulation is the ratio between the collective utility if all agents play sincerely and the collective utility if agent A plays strategically and all the other ones play sincerely. In the following, we will focus on the two agents case and Borda scoring function [3, 2], where the utility of the ith best object for an agent is $m - i + 1$.

Proposition 8 *For each $(\rhd_A, \rhd_B)$, π, we have:*

$$PM_{P,\pi,g_{Borda}}(\sigma_A) \geq 1 - \frac{2\sum_{s \in \{ps(\pi,B)_1,\dots\}} PS(s) - 2}{m^2 + m - 2PS(m)^2 + 2mPS(m) + 2PS(m)},$$

where $PS(s)$ is the number of picking stages of A until step s.

[6] The missing proofs can be found in the long version of the paper: http://recherche.noiraudes.net/resources/papers/ECAI14-full.pdf.

Proof: Let σ_A be a successful strategy for A, and u_A, u_B (resp. u'_A, u'_B) be the utilities obtained by A and B if they play sincerely (resp. A plays according to σ_A and B plays sincerely). At its i^{th} picking stage $ps(\pi, B)_i$, B can obtain in the best case its i^{th} object, and obtains in the worst case its $(i + PS(ps(\pi, B)_i)^{\text{th}}$ object. Hence $u'_B \geq u_B - \sum_{s \in \{ps(\pi,B)_1,\dots\}} PS(s)$. Moreover, since σ_A is successful, $u'_A \geq u_A + 1$. And finally, since in the best case, each agent receives his most preferred objects, we have $u_A + u_B \leq \sum_{k=1}^{PS(m)}(m - k + 1) + \sum_{k=1}^{m-PS(m)}(m - k + 1) = 1/2 \times (m^2 + m - 2PS(m)^2 + 2mPS(m) + 2PS(m))$. Hence:

$$\frac{u'_A + u'_B}{u_A + u_B} \geq 1 - \frac{\sum_{s \in \{ps(\pi,B)_1,\dots\}} PS(s) - 1}{u_A + u_B}.$$

Replacing $u_A + u_B$ by its upper bound completes the result. ■

Corollary 3 *If π is the alternating sequence (for an even number of objects),*

$$PM_{(\rhd_A,\rhd_B),ABABAB\dots,g_{Borda}}(\sigma_A) \geq 1 - \frac{m^2 + m - 4}{3m^2 + 4m}.$$

Thus, at least in this simple case, manipulation by a single agent does not have a dramatic effect on the social welfare, as it will cause only approximately 33% loss of utility in the worst case. (We also have results about the *additive* price of manipulation, that is, the worst-case difference between social welfare when A plays a sincere strategy and the social welfare when A plays strategically; due to the lack of space, we omit them.)

6 CONCLUSION

We have studied the computation of picking sequence manipulation. In the case of a single manipulator, we have found that for any number of non-manipulators and any picking sequence, finding an optimal manipulation is easy if the manipulator has additive preferences, and NP-hard in the general case. Next, finding a coalitional manipulation is easy if post-allocation object trading and side payments are allowed, and NP-hard in the other cases. Finally, we have shown that in simple cases, the price of manipulation is not significantly high.

REFERENCES

[1] S. Bouveret and J. Lang, 'Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity', *Journal of Artificial Intelligence Research (JAIR)*, **32**, 525–564, (2008).
[2] S. Bouveret and J. Lang, 'A general elicitation-free protocol for allocating indivisible goods', in *Proc. IJCAI'11*, (2011).
[3] S. Brams and D. King, 'Efficient fair division: Help the worst off or avoid envy ?', *Rationality and Society*, **17**, 387–421, (2005).
[4] S. Brams and A. Taylor, *The Win-win Solution. Guaranteeing Fair Shares to Everybody*, W. W. Norton & Company, 2000.
[5] 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, (August 2012).
[6] P. Faliszewski and A. Procaccia, 'AI's war on manipulation: Are we winning?', *AI Magazine*, **31**(4), 53–64, (2010).
[7] T. Kalinowski, N. Narodytska, and T. Walsh, 'A social welfare optimal sequential allocation procedure', in *IJCAI*, (2013).
[8] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia, 'Strategic behavior when allocating indivisible goods sequentially', in *AAAI*, (2013).
[9] D. Kohler and R. Chandrasekaran, 'A class of sequential games', *Operations Research*, **19**(2), 270–277, (1971).
[10] L.-G. Svensson, 'Strategy-proof allocation of indivisible goods', *Social Choice and Welfare*, **16**, (1999).