Manipulations in Two-Agent Sequential Allocation with Random Sequences

Yuto Tominaga, Taiki Todo, and Makoto Yokoo
Kyushu University, Motooka 744, Fukuoka, Japan
{ tominaga@agent., todo@, yokoo@ } inf.kyushu-u.ac.jp

ABSTRACT

Sequential allocation is one of the most fundamental models for allocating indivisible items to agents in a decentralized manner, in which agents sequentially pick their favorite items among the remainder based on a pre-defined priority ordering of agents (a sequence). In recent years, algorithmic issues about agents' manipulations have also been investigated, such as the computational complexity of verifying whether a given bundle of items is achievable and maximizing one's utility under a given additive utility function. In this paper we consider a slightly modified model, where the selection process is divided into rounds, each agent obtains exactly one item in each round, and the sequence per round is determined uniformly at random. It is natural to expect that finding a profitable manipulation is difficult even for the case of two agents, since a manipulator must consider exponentially many possible sequences with respect to the number of rounds due to randomization. To our surprise, however, an optimal manipulation can be computed without any exploration for exponentially decaying utilities. Furthermore, for general additive utilities, although some exploration is required, it can still be done in polynomial time with respect to the number of rounds.

CCS Concepts

•**Computing methodologies** → **Multi-agent systems;**
•**Applied computing** → *Economics;*

Keywords

Distributed Resource Allocation; Item Picking; Sequential Allocation; Uncertainty

1. INTRODUCTION

Sequential allocation, which is one of the most common research topics in the field of social science and economics [12], simply models the situation of allocating indivisible items to agents in a decentralized manner. In a traditional sequential allocation model, there is a set of indivisible items, a set of agents, and a fixed priority ordering (known as a *sequence*) over the set of agents, according to which they pick items one-by-one. Considering fairness among agents, the sequences used in practice are commonly 'balanced,' such that at every moment of the selection process the difference of the numbers of items already selected is bounded by one. This traditional model can represent several practical situations, including the draft system in many worldwide professional sports associations, such as the National Basketball Association in the US.

Sequential allocation has also been attracting a huge amount of attention from researchers in the fields of multi-agent systems, algorithmic game theory, and computational social choice, and several research issues related to these fields have been investigated. Bouveret and Lang [4] considers the problem of finding the optimal sequence in which social welfare, i.e., the sum of agents' utilities, is maximized, which was followed and extended by Kalinowski et al. [10]. Furthermore, Kalinowski et al. [11] addresses the problem from the viewpoint of sequential form games and studies the subgame-perfect equilibrium in a game defined by sequential allocation. Aziz et al. [2] investigates the computational complexity of verifying whether a given bundle is possible or necessary.

From the perspective of an agent who is willing to maximize her own utility, a natural algorithmic question asks whether she can obtain a specific bundle of items, as initiated by Bouveret and Lang [5]. Assuming that the utility of such an agent (a *manipulator*) is additive and that every other agent adopts the 'sincere' picking strategy, i.e., picking her most favorite item among the remainder at every moment of her action, they show that maximizing the manipulator's utility is solvable in polynomial time with respect to the number of items. They also conclude that the effect of such optimal manipulation on social welfare, specifically the loss of social welfare by manipulation, is not very large.

Although all of these research issues address sequential allocation with fixed sequences, some practical situations related to sequential allocation do not have such fixed sequences, including the draft system in Japan's professional baseball league. For such situations, it is natural to consider random tie-breaking in the sequential allocation problem, which is the issue we study in this paper. One of the simplest processes with random tie-breaking can be described as follows. In each round, each agent chooses one item. If there is a conflict among agents' choices, then a lottery is used to decide which agent among those having the conflict obtains the item. All of those who cannot obtain the item choose other items among the remaining ones, and the same process is repeated until every agent obtains exactly one in

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

the round. Then go to the next round and repeat the process until every item is given to an agent.

For simplicity, in this paper we consider a slightly different model and focus on the case of two agents (nevertheless, our results can still be applied to the above original model with two agents, as we discuss in Section 6). In our model, there exist $2m$ items and two agents, A and B, where m is a positive integer. The selection process is divided into m rounds, in each of which one agent obtains one item. At the beginning of each round, a coin flip determines which agent picks first; the sequence in a round is chosen uniformly at random between AB and BA, where AB means that agent A picks first followed by agent B. Here, an agent, say A, faces the problem of choosing between 'risk-taking' and a 'sure-thing.' Consider four items $\{1, 2, 3, 4\}$, where the indices indicate agent A's preference so that she prefers the lower-indexed one, agent B has preference (binary relation over the set of items) $\rhd_B$ such that $2 \rhd_B 3 \rhd_B 1 \rhd_B 4$ and plays a 'sincere' picking strategy, and the sequence in the first round is chosen as AB. By picking 2 first and deferring 1 to the next round, agent A takes a risk of losing 1 depending on the coin flip in the next round. On the other hand, by picking 1 first, she is guaranteed to obtain 1 but cannot obtain 2.

We focus on three specific computation problems, Bundle-Possibility, Bundle-Certainty and Max-Expected-Utility. Bundle-Possibility (respectively, Bundle-Certainty) defines the computation problem of verifying whether agent A can obtain a given bundle under some (resp. any) realizable sequence. On the other hand, Max-Expected-Utility defines the computation problem of finding the maximum value of her expected utility, where expectation is taken for the randomized realization of sequences. It is natural to expect that Max-Expected-Utility is not solvable in polynomial time, i.e., finding a profitable manipulation in the model with random sequences becomes much harder than in the model with deterministic sequences, since a manipulator must consider exponentially many possible sequences with respect to the number of rounds due to randomization. To our surprise, however, an optimal manipulation can be computed without any exploration in the exponentially decaying utility case. Furthermore, even in the general additive utility case, although some exploration is required to find an optimal manipulation, it can be done in polynomial time with respect to the number of rounds.

Let us introduce how we can interpret our technical results. From the viewpoint of a mechanism designer who organizes a sequential allocation with random sequences, those results have negative implications. If an agent wants to maximize her utility, finding the best manipulation is easy, and she will not use the sincere picking strategy. This means that introducing randomization in tie-breaking does not help the mechanism designer provide agents incentive to behave sincerely, at least for the case of two agents. However, our analysis is just a first step; considering more than two agents in the model and discussing other criteria such as social welfare and fairness remain an open and interesting future direction.

The paper is organized as follows. Section 2 compares our model and its contributions with previous works. Section 3 introduces our model and explains several terms used in the technical parts. Section 4 defines four decision problems, Bundle-Possibility, Bundle-Certainty, Find-Best-Possible and Find-Best-Certain, and shows that all can be solved in polynomial time. Section 5 defines the Max-Expected-Utility problem and shows that it can also be solved in polynomial time. Particularly, for the case of exponentially decaying utilities, we show that the problem can be solved much faster, without any exploration. Section 6 shows an application of our results to the model with random tie-breaking. Section 7 discusses related existing works and raises possible future research questions. Finally, Section 8 concludes the paper.

2. RELATED WORKS

Some prior works have dealt with the decentralized allocation of indivisible items with randomization. The contested pile methods, proposed by Vetschera and Kilgour [14] and followed by Brams et al. [6], create a decentralized allocation model, in which a conflict over an item is solved by putting it into a place called a *contested pile* and deciding its allocation after the picking sequence is completed. The parallel elicitation-free protocol, proposed by Huang et al. [9] is another class of decentralized procedures, in which every agent has the same probability to get each item. Both randomizations in these models differ from ours, and thus their results have no direct implication for our model.

In the literature of social choice and mechanism design, some researches also deal with the randomized allocation of indivisible items and focus on centralized allocation mechanisms. Budish and Cantillon [7] analyzes the random serial dictatorship mechanism, which is a well-known type of scheme for allocating indivisible items in a centralized manner. Bogomolnaia and Moulin [3] analyzes the probabilistic serial rule, which is another centralized mechanism that will be explained in more detail in Section 7. Analyzing manipulation complexity is also common in the fields of artificial intelligence and multi-agent systems, even for the problem of allocating indivisible items [1, 8, 13].

3. MODEL

We first introduce the model of sequential allocation with random sequence dealt with in this paper. There is a set $\mathcal{O}$ of $2m$ items, which are indivisible and distinct, and a set of two agents, A and B.

The sequential allocation we consider is described as follows. In each *round*, the rule flips a fair coin to determine which agent picks an item first. If 'heads' appears, the sequence in the round (or the *partial sequence*) is set to AB, where agent A first picks an item among the remaining ones, followed by agent B. If 'tails' appears, the sequence is set to BA, where agent B first picks, followed by agent A. Then go to the next round and repeat this process as long as items remain. Note that by definition, there are exactly m rounds and each agent has exactly one turn to pick an item in each round, and thus both agents finally obtain m items. Let $\pi \in \Pi_{\mathcal{B}} = \{\mathrm{AB}, \mathrm{BA}\}^m$ indicate a sequence of the realizations of 'heads' and 'tails' by coin flips in each of m rounds. Here $\Pi_{\mathcal{B}}$ corresponds to the set of possible 'recursively balanced' sequences [2].

Next let us introduce several terms that will be used to define the problems investigated in this paper. Consider agent A a manipulator, assuming that B plays a 'deterministic, simple picking' strategy [5], which is referred to as a *sincere strategy* in this paper. Let $\rhd_B$ indicate a strict ordering of all items $\mathcal{O}$, representing the *preference* of agent

B. Agent B picks, in each of her turns, her most favorite item among those remaining, and agent A knows this fact. Let $u_A : \mathcal{O} \to \mathbb{R}_+$ be a utility function of agent A. We assume, in this entire paper, that agent A's utility is *additive*, meaning that for given utility function u_A of agent A and bundle $X \subseteq \mathcal{O}$, $u_A(X) := \sum_{i \in X} u_A(i)$. We usually describe the items by indices $1, \ldots, 2m$ based on agent A's preference so that A prefers a lower-indexed one, i.e., for any $i, j \in \mathcal{O}$, $u_A(i) > u_A(j)$ if and only if $i < j$. We also say that utility function u_A is *exponentially decaying* if $u_A(i) > u_A(j) \to u_A(i) \geq 2u_A(j)$ for every $i, j \in \mathcal{O}$ Sometimes we refer to such a exponentially decaying utility function as v. Furthermore, for given utility function u_A and probability distribution function F associated with mass function f over $2^{\mathcal{O}}$, expected utility $u_A(F)$ of agent A is defined as $\mathbb{E}_{X \sim F}[u_A(X)] = \sum_{X \in 2^{\mathcal{O}}} u_A(X) \cdot f(X)$. We write $\rhd_B$ and u_A as $\rhd$ and u, respectively, if they are clear from the context.

Now we are ready to introduce some additional terms and notions related to agent A's strategy in our model of sequential allocation. For given bundle X of items and preference $\rhd$, let $L_{X,\rhd}$ be the list of the items in X that are ordered based on $\rhd$. Particularly, for given $\rhd$, let L^{init} be the sorted list of all the items in $\mathcal{O}$, where $L^{init} := L_{\mathcal{O},\rhd}$, which we sometimes denote as $L^{init} = (o_1, o_2, \ldots, o_{2m-1}, o_{2m})$. We omit both subscripts X and $\rhd$ for simplicity and write L whenever they are clear from the context. Furthermore, for notation simplicity, let $u(L_{X,\rhd}) := u(X)$.

A (standard) *plan* δ with length $k \leq m$ specifies, for given list L of remaining items and partial sequence $s \in \{\text{AB}, \text{BA}\}$ realized by a coin flip in the current round, a sequence of k choices for subsequent k rounds. Plan δ is represented as a binary tree, where each node corresponds to an item, the root node indicates the item that agent A is going to pick in the current round, and the left (resp. right) child indicates the item that A will pick when the partial sequence in the next round is AB (resp. BA). We say plan δ with length k is *feasible* for given list L of the remaining items and partial sequence s in the current round if for any possible sequence of k partial sequences that will be realized by coin flips in subsequent k rounds, agent A can pick the item indicated by the corresponding node. In other words, the item remains at the turn of A's choice. A *secure plan* is a special form of plans, which indicates, for each round, the selection of the same item regardless of the partial sequence. Hence, a secure plan can also be represented as a list, where the first (second, and so on) element of the list indicates the item A picks in the first (second, and so on) round.

Although a plan has a sufficient power of representation, it has $O(2^k)$ nodes that cannot be efficiently described. Therefore we introduce a compact representation of the plans called a *strategy*. Strategy σ is a function that returns an item, $\sigma(s, L)$, that agent A should pick for given partial sequence s and list L of the remaining items with even length. For given strategy σ and list L, let $E(L, \sigma)$ represent agent A's expected utility when she uses strategy σ for list L of items. Assume that $E(L, \sigma) = 0$ whenever L is empty list (). Analogously, let $E_{\text{AB}}(L, \sigma)$ (resp. $E_{\text{BA}}(L, \sigma)$) represent A's expected utility under the condition that the partial sequence in the current round is AB (resp. BA), i.e., $E_{\text{AB}}(L, \sigma) := u(i) + E(L', \sigma)$, where i is the item that σ indicates to pick at the current round under sequence AB, i.e., $i = \sigma(\text{AB}, L)$ and L' is the list of remaining items at the beginning of the next round. By the assumption that we focus on a fair coin, $E(L, \sigma) = 0.5 \cdot E_{\text{AB}}(L, \sigma) + 0.5 \cdot E_{\text{BA}}(L, \sigma)$. Let Σ be the set of all strategies of agent A for the problem. Furthermore, let $E^*(L)$ be A's optimal expected utility achieved using one of the optimal strategies for L. Formally, $E^*(L) = \max_{\sigma \in \Sigma} E(L, \sigma)$. As is the case with E, let $E^*_{\text{AB}}(L)$ (resp. $E^*_{\text{BA}}(L)$) represent A's optimal expected utility under the condition that the partial sequence in the current round is AB (resp. BA). By definition, $E^*(L) = 0.5 \cdot E^*_{\text{AB}}(L) + 0.5 \cdot E^*_{\text{BA}}(L)$.

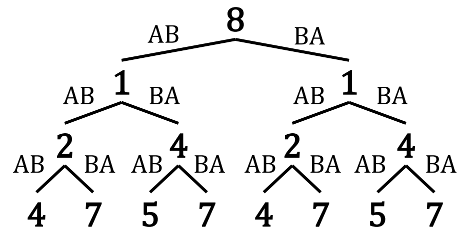

Figure 1: A plan is represented as a binary tree.

The following example explains how these terms work for a given problem instance.

Example 1. *Consider two agents, A and B, and eight items $\mathcal{O} = \{1, 2, \cdots, 8\}$ (i.e., $m = 4$). Assume that B has preference $\rhd$ such that $3 \rhd 8 \rhd 6 \rhd 2 \rhd 1 \rhd 5 \rhd 4 \rhd 7$ over them. Furthermore, we consider two utility functions, u and v, that can be held by agent A. We assume that $u(i) = 9 - i$ for any $i \in \mathcal{O}$, i.e., u is the Borda scoring function, and $v(i) = 2^{8-i}$ for any $i \in \mathcal{O}$, i.e., v is exponentially decaying.*

Plan δ with length 4 of agent A is presented as a binary tree in Fig. 1. This plan is feasible for $L^{init} = (3, 8, 6, 2, 1, 5, 4, 7)$ and both $s = AB$ and $s = BA$. However, it is not secure, because for some list L of the remaining items, say $L = (2, 5, 4, 7)$ at the rightmost node with label 1, it chooses a different item based on the realized partial subsequence, explaining why this plan cannot be represented as a list.

As a compact representation, we define a strategy σ based on this plan δ. Strategy σ is a function such that $\sigma(AB, (2, 5, 4, 7)) = 2$, $\sigma(BA, (2, 5, 4, 7)) = 4$, etc. Under above utility function u, expected utility $E(L^{init}, \sigma)$ achieved by strategy σ is 18.25. Similarly, under v, the expected utility achieved by the strategy is 176.

4. POSSIBLE AND CERTAIN BUNDLES

In this section we consider the computation problem of verifying whether agent A can obtain a given bundle. As an analogy from the possible and necessary winner problems in the literature of computational social choice, we define two variants of the problem for our model of sequential allocation with random sequences.

Definition 1 (Best Possible (Certain) Bundle). *For given $\rhd$ and current partial sequence s, bundle $X \subset \mathcal{O}$ is said to be* possible *(resp.* certain*) under $\rhd$ and s if for some (resp. any) sequence $\pi \in \Pi_{\mathcal{B}}$ of realizations that begins with s, there exists feasible plan δ that achieves X. Furthermore, for given $\rhd$, s, and u, the* best possible *(resp.* best certain*) bundle is the possible (resp. certain) bundle X that maximizes $u(X)$ among all the possible (resp. certain) bundles under $\rhd$ and s.*

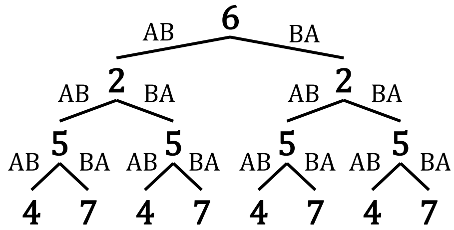

Figure 2: A feasible plan that achieves $(6, 2, 5, 4)$ under the most optimistic sequence.

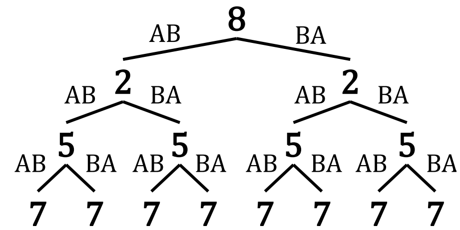

Figure 3: A feasible secure plan that achieves $(8, 2, 5, 7)$ under the most pessimistic sequence.

Definition 2 (Bundle-Possibility (-Certainty)). *For given $\rhd$, s, and $X \subset \mathcal{O}$, is X possible (resp. certain) under $\rhd$ and s ?*

The following theorem shows that both problems can be efficiently solved by agent A, which has a negative implication that agent A will behave insincerely even if her computation power is bounded. These results can be considered as a generalization of the results in Bouveret and Lang [5].

Theorem 1. *For any given $\rhd$, s, and $X \subset \mathcal{O}$, both* Bundle-Possibility *and* Bundle-Certainty *can be solved in polynomial time with respect to m.*

Proof. By Lemma 4 (resp. Lemma 5) that we will provide in Appendix, it suffices to verify whether X can be achieved under $\rhd$ and specific realization $\bar{\pi} = s\text{ABAB}\ldots\text{AB}$ (resp. $\underline{\pi} = s\text{BABA}\ldots\text{BA}$). Bouveret and Lang [4] proved that for any given $\rhd$ and any given sequence π, such a verification can be solved in polynomial time with respect to the number of items, $2m$ in our model. Therefore it is also solved in polynomial time with respect to m. □

The next problem of finding the best possible (resp. certain) bundle is also related to the above problem. Actually it can also be solved efficiently.

Definition 3 (Find-Best-Possible(-Certain)). *For given $\rhd$, s, and u, return the best possible (resp. best certain) bundle $X \subset \mathcal{O}$ under $\rhd$, s and u.*

Theorem 2. *For any given $\rhd$, s, and u, both* Find-Best-Possible *and* Find-Best-Certain *can be solved in polynomial time with respect to m.*

Proof. By Lemma 4 (resp. Lemma 5), it suffices to return bundle X that maximizes $u(X)$ among all the bundles that can be achieved under $\rhd, \bar{\pi}$ (resp. $\underline{\pi}$). Bouveret and Lang [5] proved that for any given $\rhd$, any given s, and any given additive utility function u, such a bundle can be uniquely found in polynomial time with respect to the number of items, which is given as $2m$ in our model. Therefore it can still be found in polynomial time with respect to m. □

Example 2. *Consider again the same problem with Example 1. For example, bundle $(6, 2, 5, 4)$ is possible under $s = AB$ since it can be achieved by a feasible plan under $\rhd$ and $\bar{\pi} = ABABAB\ldots$, e.g., by the one described in Fig. 2, while it is not certain since it cannot be achieved by any feasible plan under $\rhd$ and $\underline{\pi} = ABBABA\ldots$. On the other hand, bundle $(8, 2, 5, 7)$ is possible and certain under $s = AB$ since it can be achieved by a feasible secure plan under $\underline{\pi}$, e.g., by the one described in Fig. 3. By a similar argument (shown in the proof of Theorem 2), we can confirm that under both utility functions, u and v, $(3, 2, 1, 4)$ and $(2, 1, 4, 7)$ are the best possible and best certain bundle under $s = AB$, respectively.*

5. MAXIMIZING EXPECTED UTILITY

Another important question from the manipulator's viewpoint is how much expected utility she can get upon the uncertainty of the realization of subsequent sequences. First we formally define the problem.

Definition 4 (Max-Expected-Utility). *For given $\rhd$, s, and u, return the maximum expected utility $E^*_s(L^{init}) = \max_{\sigma\in\Sigma} E_s(L^{init}, \sigma)$.*

To describe the main theorems in this section, we introduce an additional term called *critical depth*, which is related to agent A's strategy and is useful for the proof. Let L be the list of remaining items and X be the best possible bundle for L and current partial sequence s. Then the critical depth is integer $d \in \mathbb{N}$ such that A first fails to pick the d-th item in $L_{X,\rhd}$ under $\underline{\pi} = s\text{BABA}\cdots\text{BA}$ by picking from the top of list $L_{X,\rhd}$. When A does not fail until the end, i.e., when X is certain under $\rhd$ and s, critical depth d is defined by $|X|+1$, i.e., the number of remaining rounds (including the current one) plus one. Letting d be the critical depth, we observe that $d \geq 2$. Furthermore, letting x_d be the d-th top item in $L_{X,\rhd}$, we observe that x_d is always at an odd-numbered position in L, because (i) under the above strategy of A and the sincere strategy of B, exactly $2d-2$ items in total that appear before x_d in L^{init} will be picked before round d, and (ii) x_d is on the top of the list of remaining items at the beginning of round d, since x_d will be picked by agent B when $s = \text{BA}$ by definition.

We introduce another technical term called an *admissible secure plan*, which is a refinement of secure plans. Let $X_d = \{i \in X \mid i \rhd x_d \vee i = x_d\}$, i.e., the set of items in X that appears before x_d (including x_d). $|X_d| = d$ clearly holds. Then an admissible secure plan is a secure plan of length $d-1$ that picks the $d-1$ items among X_d in the first $d-1$ rounds in the order specified by $\rhd$. Thus, d admissible secure plans usually exist. When d is defined by the number of remaining rounds plus one, then the admissible secure plan just picks the items of X_d in the order specified by $\rhd$. In this case, only one admissible secure plan exists. Note that any admissible secure plan is feasible, because any bundle Z that satisfies $Z \subset X_d, |Z| = d-1$ is certain under s and $\rhd$.

Also, for given utility u, we call the admissible secure plan, which takes $d-1$ items among X_d except for $j' =$

$\arg\min_{j \in X_d} u(i)$, i.e., the least preferred item, a *prudent admissible secure plan.*

5.1 Exponentially Decaying Utilities

We first focus on a specific form of utility functions and show that the problem can be solved quite easily, in the sense that we do not need any exploration technique. More specifically, when the utility of agent A is restricted to exponentially decaying ones, we can solve the maximization problem without any exploration or comparison of the expected utilities during the whole search process.

THEOREM 3. *For any given $\rhd$, s, and A's utility function v that is exponentially decaying,* MAX-EXPECTED-UTILITY *can be solved in polynomial time with respect to m.*

PROOF SKETCH. Let us show the optimality of the following strategy $\widehat{\sigma} \in \Sigma$ of agent A: pick the item specified by the prudent admissible secure plan. Such strategy $\widehat{\sigma}$ is obviously feasible and also runs in polynomial time because of its 'greedy' nature, i.e., it just follow a particular secure plan and does not compare expected utilities that may happen in the future from the current state. Thus, $E_s(L, \widehat{\sigma})$ can be computed in polynomial time.

Now we show an intuition of why $\widehat{\sigma}$ achieves the optimal expected utility. Let v be the exponentially decaying utility function. Here, assume for the sake of contradiction that another strategy σ' achieves an optimal expected utility that is strictly better than that of $\widehat{\sigma}$. Then, during the first $d-1$ rounds in a possible partial sequence, it picks either (a) $d-1$ items among X_d that are different from those that $\widehat{\sigma}$ picks or (b) $d-1$ items some of which does not belong to X_d; otherwise σ' coincides with $\widehat{\sigma}$ by applying the same argument until the end.

For case (a), σ' assuredly obtains item $j' = \arg\min_{j \in X_d} v(i)$, by taking a risk that she cannot obtain another item that she prefers to j' among X_d with probability that exceeds 1/2. Since v is exponentially decaying, the expected utility of such strategy σ' cannot dominate that of $\widehat{\sigma}$.

Case (b) is more complicated. σ' obtains bundle $Y \not\subseteq X_d$, which satisfies $|Y| = d-1$ and is possible under s and $\rhd$, during the first $d-1$ rounds. In this case, we can construct another strategy that dominates σ' by obtaining $d-1$ items from X_d during the first $d-1$ rounds and then assuredly obtaining those items $(Y \cap X) \setminus X_d$ during the later rounds, violating the assumption of the optimality of σ'. □

Theorem 3 is part of Theorem 4, which we present in the next subsection. However, the following example demonstrates that the optimal strategy for the case of exponentially decaying utilities runs without any exploration and thus is much faster than a strategy that can be applied to general additive utilities.

EXAMPLE 3. *Consider again the same problem with Example 1, and focus on utility function v that is exponentially decaying. When $s = AB$ in the first round, the (sorted) best possible bundle is $(3, 2, 1, 4)$ and the critical depth is 3. Then the strategy mentioned in the above proof suggests picking 2 in the first round (and 1 in the second round).*

At the beginning of the third round, the list of remaining items is $(6, 5, 4, 7)$. In this round, the best possible bundle is $(5, 4)$ and the critical depth is 2, regardless of s. Then the strategy suggests picking 4. In the final round, the strategy suggests picking the best remaining item. Thus, the strategy absolutaly takes $1, 2, 4$, and 5 or 7 with probability 0.5. The expected utility is 213.

When $s = BA$ in the first round, the plan suggests picking exactly the same items in the case where $s = AB$. Thus, the expected utility is 213.

5.2 General Additive Utilities

Unlike the case of exponentially decaying utility functions, the maximization problem seems much more difficult when utility functions are general. An intuition of the difficulty is that, to calculate the maximized expected utility, agent A needs to explore its future by comparing all admissible secure plans, appropriately choose one, and repeat this procedure until the best possible bundle at the moment becomes certain. However, the following theorem shows that the problem can still be solved in polynomial time. One main reason is that, although the whole procedure looks complicated, the number of states at which she needs to explore its future is bounded by a polynomial with respect to the number of items $2m$.

The following lemma implies that we can restrict our attention to plans/strategies that are based on admissible secure plans without losing optimality. Note that this is a generalization of case (b) in the proof of Theorem 3 for general additive utility functions.

LEMMA 1. *For any L of even length and s, there exists admissible secure plan L' such that $E^*_s(L) = u(L') + E^*(L'')$ holds, where L'' is the list of remaining items after executing L'.*

PROOF SKETCH. Let X be the best possible bundle under s and $\rhd$, let d be the critical depth under the current partial sequence s and the ordering $\rhd$, $X_d := \{i \in X \mid i \rhd x_d \vee i = x_d\}$, and let Y be an arbitrary bundle that is possible under s and $\rhd$ and satisfies both $|Y| = d-1$ and $Y \not\subset X_d$. Then we can construct $Z \subset X_d$ that satisfies $|Z| = d-1$ and $u(Z) + E^*(L'_Z) \geq u(Y) + E^*(L'_Y)$, where L'_Z (or L'_Y) is the list of remaining items after A takes items in Z (or Y) and B takes items based on $\rhd$ from L.

Now let Z^* be the optimal bundle among those Z that satisfies the above conditions. Such an optimal Z^* is certain under s and $\rhd$, and thus the optimal strategy must pick Z^* during the first $d-1$ rounds in the order specified by $\rhd$. This coincides with a strategy based on the admissible secure plan represented by $L_{Z^*,\rhd}$, which concludes the proof. □

THEOREM 4. *For any given $\rhd$, s, and u,* MAX-EXPECTED-UTILITY *can be solved in polynomial time with respect to m.*

To prove the statement, we first introduce the following specific form of schemes that we call *strategies based on admissible secure plans* (SASP). This is obviously a generalization of the strategy we mentioned in the proof of Theorem 3.

DEFINITION 5 (SASP). *The* strategy based on admissible secure plans $\sigma^*(s, L)$ *is defined as follows:*

1. *Enumerate all admissible secure plans.*
2. *For each admissible secure plan L', let L'' denote the list of remaining items after executing L'. If $E(L'', \sigma^*)$ was not calculated before, then recursively run*

$\sigma^*(AB, L'')$ and $\sigma^*(BA, L'')$ to obtain $E_{AB}(L'', \sigma^*)$ and $E_{BA}(L'', \sigma^*)$, respectively. Then set $E(L'', \sigma^*)$ to their average.

3. *Choose L' that maximizes $u(L')+E(L'', \sigma^*)$ and return the first item in L' and $u(L') + E(L'', \sigma^*)$.*

In the following we prove Theorem 4 by separately showing that (i) SASP is optimal (in Proposition 1) and that (ii) SASP runs in polynomial time in m (in Proposition 2).

PROPOSITION 1 (OPTIMALITY). *For any L, $E(L, \sigma^*) = E^*(L)$ holds. Hence, σ^* is an optimal strategy.*

PROOF. We prove the statement by mathematical induction with respect to the length of list L. Note again that the length of L is always even since we are focusing on recursively balanced sequences. For the base case, in which $L = ()$ holds, the statement is obviously true by definition. Now let us move to the induction step.

Assuming that $E(L, \sigma^*) = E^*(L)$ holds for any list L s.t. $|L| \leq 2n$ (where $n = 0, 1, \ldots$), we consider arbitrary list L s.t. $|L| = 2n+2$. First we focus on the case where the partial sequence in the current round is AB.

If the critical depth is $n + 2$, i.e., if the best possible bundle X can be achieved even under subsequence ABBABA...BA, then there exist a unique admissible secure plan that achieves the best possible bundle, and thus it holds that $E_{\text{AB}}(L, \sigma^*) = u(L_{X,\triangleright}) + E((), \sigma^*) = E^*(L)$.

On the other hand, if the critical depth does not equal $n+2$, then from Lemma 1, there exists admissible secure plan L' such that $E^*_{AB}(L) = u(L') + E^*(L'')$, where L'' is the list of remaining items after executing L'. Since it is the case that $|L''| \leq 2n$, by the assumption, $E^*(L'') = E(L'', \sigma^*)$ holds.

Therefore, it holds that $E_{\text{AB}}(L, \sigma^*) \geq E^*_{\text{AB}}(L)$. By definition of E^*_{AB}, this implies that $E_{\text{AB}}(L, \sigma^*) = E^*_{\text{AB}}(L)$. By a similar argument, it also holds that $E_{\text{BA}}(L, \sigma^*) = E^*_{\text{BA}}(L)$. Thus, $E(L, \sigma^*) = E^*(L)$ holds. □

We then show that SASP can be calculated in polynomial time. First, we introduce an additional concept of *reachability* for bundles. Let $\mathcal{L}$ be a series of sorted lists defined below, each of which can be obtained by eliminating all but one element in the first $2t-1$ elements, for each $t \in 1, \ldots, m$:

$$\begin{aligned}\mathcal{L} = \{&(o_1, o_2, \ldots, o_{2m}),\\ &(o_1, o_4, \ldots, o_{2m}), (o_2, o_4, \ldots, o_{2m}), (o_3, o_4, \ldots, o_{2m}),\\ &(o_1, o_6, \ldots, o_{2m}), (o_2, o_6, \ldots, o_{2m}), \ldots, (o_5, o_6, \ldots, o_{2m}),\\ &\vdots\\ &(o_1, o_{2t}, \ldots, o_{2m}), \ldots, (o_{2t-1}, o_{2t}, \ldots, o_{2m}),\\ &\vdots\\ &(o_1, o_{2m}), (o_2, o_{2m}), \ldots, (o_{2m-1}, o_{2m})\}.\end{aligned}$$

Note that $|\mathcal{L}| = 1 + 3 + \cdots + (2m - 1) = m^2$. List L is said to be *reachable* from L^{init} (by admissible secure plans) if there exists a series of partial sequences and an admissible secure plan such that after the execution of the admissible secure plan for the series of partial sequences, the list of remaining items becomes L.

The following lemma helps us to show the polynomial time computability, which implies that the number of lists of remaining items possible after any sequence of applications of an admissible secure plan is at most $|\mathcal{L}| = m^2$.

LEMMA 2. *If L is reachable from L^{init}, $L \in \mathcal{L} \cup \{()\}$ holds.*

PROOF. We prove the statement by mathematical induction with respect to k the number of times that an admissible secure plan is applied. For the base case in which $k = 0$, such a reachable list that is obtained without applying an admissible secure plan must be L^{init}, which is in set $\mathcal{L}$ by its definition.

Now we show the induction step. Assuming that any list that is reachable by applying an admissible secure plan k' times is in set $\mathcal{L}$, consider arbitrary list L that is reachable by applying an admissible secure plan $k' + 1$ times. At the moment of the $k' + 1$-st application of an admissible secure plan, remaining list $\widehat{L}$ is in $\mathcal{L}$, because of the fact that it is reachable by applying an admissible secure plan k' times and the assumption.

Here let d and L' be the critical depth and the sorted list of the items in the best possible bundle at the moment, respectively. If $d = |L'| + 1$, i.e., if all the remaining items in the best possible bundles can be obtained, then L is $()$. On the other hand, if it is not the case, then such list L is obtained by eliminating the top $2d - 2$ items among the top $2d - 1$ items in $\widehat{L}$ by the definition of the admissible secure plan and the sincere strategy of B, which results in a list belonging to $\mathcal{L}$. □

PROPOSITION 2 (POLY-TIME). *For any L that is reachable from L^{init} and any $s \in \{AB, BA\}$, $\sigma^*(s, L)$ is calculated in polynomial time with respect to m.*

PROOF. Strategy σ^* is applied only for each list that is reachable from L^{init}, which takes at most m^2 applications by Lemma 2. In each application of σ^*, the number of admissible secure plans is $O(m)$. Any admissible secure plan just needs operations (e.g., calculating the best possible bundle) that only take polynomial runtime. As a result, strategy σ^* can be calculated in polynomial time with respect to m. □

EXAMPLE 4. *Consider again the same problem with Example 1 with Borda scoring function u. When $s = AB$ in the first round, the (sorted) best possible bundle is $(3, 2, 1, 4)$, and the critical depth is 3. Then strategy σ^* considers obtaining two out of the three top items in $(3, 2, 1, 4)$ without taking any risk.*

Obtain $(3, 1)$**:** *After taking them, the remaining items at the beginning of the third round will be $(2, 5, 4, 7)$. When $s = AB$ in the third round, the best possible bundle is $(2, 4)$, and the critical depth is 2. Then the strategy next considers obtaining one out of the two items $(2, 4)$ without taking any risk.*

- **Obtain** (2)**:** *After taking item 2, the remaining items at the beginning of the fourth round will be $(4, 7)$. When $s = AB$ in the fourth round, the best possible bundle is (4), which is certain. When $s = BA$, the best possible bundle is (7), which is also certain. As a result, strategy σ^* returns an expected utility of $(u(4) + u(7))/2$.*
- **Obtain** (4)**:** *After taking item 4, the remaining items at the beginning of the fourth round will be $(5, 7)$. By a similar argument, it returns an expected utility of $(u(5) + u(7))/2$.*

Comparing $u(2) + (u(4) + u(7))/2$ and $u(4) + (u(5) + u(7))/2$, it suggests picking 2 when $s = AB$. By a *similar argument, it suggests picking* 4 *when* $s = BA$*. As a result, it returns an expected utility of*

$$u(3) + u(1) + \frac{(u(2) + \frac{u(4)+u(7)}{2}) + (u(4) + \frac{u(5)+u(7)}{2})}{2}.$$

Obtain $(3, 2)$**:** *By a similar argument, it returns*

$$u(3) + u(2) + \frac{(u(1) + \frac{u(4)+u(7)}{2}) + (u(4) + \frac{u(5)+u(7)}{2})}{2}.$$

Obtain $(2, 1)$**:** *it returns*

$$u(2) + u(1) + u(4) + \frac{u(5) + u(7)}{2}.$$

Comparing the expected utilities based on u*, it suggests picking* 3 *in the first round when* $s = AB$*. By a similar argument, it suggests picking* 2 *in the first round when* $s = BA$*.*

6. APPLYING TO RANDOM TIE-BREAKING

As an application of the results presented in the previous sections, we now show the original model of sequential allocation with random tie-breaking that was already mentioned in the introduction. Since the discussion in the previous sections is focused on the case of two agents, here we continune to focus on the case of two agents: A and B.

The process of sequential allocation with random tie-breaking is described as follows. In each round, each agent bids on an item that still exists. If there is no conflict, each gets the item on which she bids. If there is a conflict, the 'winner' is decided by flipping a fair coin. The other agent picks another item that is still present. The process is repeated until all the items are picked.

In this model we also assume that agent B just uses the sincere picking strategy based on her preference $\rhd$. Although the model with random tie-breaking is slightly different from random sequences, some strategies in the random sequence model remain valid for this random tie-breaking. More precisely, there exists a bijection between the set of all possible strategies in the model with random tie-breaking and a certain set of strategies in the random sequence model, such that corresponding strategies give the same expected utility for the random sequence model and the random tie-breaking model. Furthermore, SASP is actually in the set. Therefore, the following statement holds.

THEOREM 5. *Given any* $\rhd$ *and* u*, the maximum expected utility in sequential allocation with random tie-breaking can be calculated in polynomial time with respect to* m*.*

Due to space limitations, we omit the formal proof. Instead, we demonstrate how SASP works for this model.

EXAMPLE 5. *Consider again the same problem with Example 1 with Borda scoring function* u*. Agent* A *first computes SASP* σ^* *to determine the item on which she first bids, by assuming the random sequence model and that* $s = AB$ *in the first round. As we already observed in Example 4,* σ^* *returns item* 3*.*

Since B *also bids on* 3 *first based on the sincere picking strategy, there will be a coin flip. If* A *wins, she gets* 3*. Then in the next round she bids on* 1 *as* σ^* *suggests. After this,* A *computes* σ^* *again for the remaining items, and so on.*

On the other hand, if A *loses in the first round, she then computes* σ^* *again by assuming that the first partial sequence in the random sequence model is* BA*, which suggests picking* 2 *in the first round,* 1 *in the second round, and* 4 *in the third round, which ensures that no conflict occurs within the three rounds. Then* A *computes* σ^* *again for the remaining items.*

7. DISCUSSIONS

One possible question about our sequential allocation model with random sequences is how it differs from the well-known probabilistic serial rule [3]. Here we show an example that results in different outcomes under these two rules with sincere strategies chosen by both agents. Consider a situation with two agents, A and B, where B's preference $\rhd$ is such that $1 \rhd 4 \rhd 2 \rhd 3$, and both A and B choose the sincere picking strategy (not alike our analysis where A seeks the optimal strategy). Under the probabilistic serial rule, both first start to 'eat' item 1, and then after finishing 1, agent A eats 2 while B eats 4. Finally, both come to item 3 at identical timing, which results in a utility of $0.5 \cdot u(1) + u(2) + 0.5 \cdot u(3)$. On the other hand, under sequential allocation with random sequence, agent A's utility is $0.5 \cdot u(1) + 0.75 \cdot u(2) + 0.75 \cdot u(3)$. Note again that this utility is under A's sincere strategy. Verifying whether they still differ in equilibrium is an open question.

Another possible question is how the coincidence result presented in Section 6 could be applied to cases with more than two agents. We observed that sequential allocations with random sequences and random tie-breaking do not coincide when there are more than two agents. To demonstrate the intuition, consider the following three agents: agent A who is willing to pick bundle $(1, 2, 3)$ in this order, agent B who picks items sincerely with her preference $\rhd_B : (1, 3, 2)$, and agent C who also picks items sincerely with her preference $\rhd_C : (3, 2, 1)$. Under our model with random sequences, it is possible that partial sequence ABC is realized, which results in an allocation such that A, B, and C obtain 1, 3, and 2, respectively. On the other hand, in sequential allocation with random tie-breaking, agent C solely bids on 3 in the first round, while both A and B bid on 1. Therefore, regardless of the result of the coin flip for tie-breaking between A and B, C obtains 3, which cannot occur in the random sequence model.

From the perspective of fair allocation, perhaps constantly flipping a fair coin is not natural. This actually makes sense; when the 'heads' side appears in the first round, it seems fairer to use a biased coin so that it is likely to be 'tails' in the next round. However, to the best of our knowledge, no such discussion exists about the optimal (fairest) probability distribution of coin flips. In this sense, our contribution is a first step toward a new research direction on optimizing random sequences with respect to some fairness criteria, such as envy-freeness and/or proportionality.

8. CONCLUSIONS

In this paper we consider a new model of sequential allocation with random sequences. We focus on the case of two agents and additive utilities of a manipulator and show that BUNDLE-POSSIBILITY, BUNDLE-CERTAINTY, FIND-BEST-POSSIBLE, FIND-BEST-CERTAIN, and MAX-EXPECTED-UTILITY can be solved in polynomial time with respect to the number m of rounds. When the manipulator's utility is further restricted to be exponentially decaying, MAX-EXPECTED-UTILITY can be solved much faster than the general case. As an application of MAX-EXPECTED-UTILITY, we show that for a different model of sequential allocation with random tie-breaking, MAX-EXPECTED-UTILITY can also be solved in polynomial time. Besides the questions raised in Section 7, several related future research issues remain. Considering more than two agents in our model is one obvious future direction. Also discussing how such a selfish manipulation affects social welfare is another, as Bouveret and Lang [5] do for a model with deterministic sequences.

Acknowledgments

This work was partially supported by JSPS KAKENHI Grants Number 24220003, 26540118, and 26730005.

We thank Haris Aziz, Jérôme Lang, and anonymous referee for their helpful comments and suggestions. All errors are our own.

APPENDIX

Here we give supplementary materials used for the theorems in Section 4. We first introduce some additional notation for the proof.

First, for standard plan δ of length k defined for partial sequence s, list L, and sequence $\pi \in s\{AB, BA\}^{k-1}$, let $\delta(\pi)$ denote the list of items taken by δ when the realized sequence is π, i.e., $\delta(\pi)$ is one particular path in the tree represented by δ. We call $\delta(\pi)$ an induced plan of δ in π. Let X denote the set of items in $\delta(\pi)$. We say $\delta(\pi)$ is *regular*, if $\delta(\pi) = L_{X,\rhd}$, i.e., the induced plan takes the items in X in the order of $\rhd$.

Let $sb(\rhd, i) = |\{j \in \mathcal{O} \mid j \rhd i\}|$ be the number of items in $\mathcal{O}$ that are strictly better than item i under given ordering $\rhd$. Also, let $blk(\rhd, X, i) = sb(\rhd, i) - |\{j \in X \mid j \rhd i\}|$ be the number of "blanks" before item i, which means that when agent A is going to obtain X, agent B does not pick i in any round before $blk(\rhd, X, i)$. When it is clear from the context, we write $blk(i)$ instead of $blk(\rhd, X, i)$. Also, given π and $r \in \{1, \ldots, m\}$, let π_r indicate the partial sequence in π for round r. Then define $\rho(\pi_r)$ such that $\rho(\pi_r) = 1$ if $\pi_r =$ AB and $\rho(\pi_r) = 0$ otherwise.

The following corollary is derived from Proposition 8 in Bouveret and Lang [5].

COROLLARY 1. *Given $\rhd$ and sequence π, X can be achieved under $\rhd$ and π if and only if there exists feasible plan δ such that its induced plan $\delta(\pi)$ is regular and achieves X under $\rhd$.*

PROPOSITION 3. *For given $\rhd$, π, and $X \subseteq \mathcal{O}$, a regular induced plan $\delta(\pi)$ achieves X under $\rhd$ if and only if $blk(x_r) \geq a(\pi, r)$ holds for any $r \in \{1, \ldots, |X|\}$, where $a(\pi, r) = r + \rho(\pi_r) - 1$, representing the number of actions (equally the number of items) B has already taken prior to A's r-th action.*

PROOF. The "if" part is obvious, because from the meaning of blk, x_r is still available for A during the first $blk(x_r)$ rounds from the current state. Therefore it suffices to show the "only if" part. Assuming that the regular induced plan $\delta(\pi)$ achieves X, we show the conclusion by mathematical induction for r.

As the base case, consider $r = 1$; we show that $blk(x_1) \geq a(\pi, 1) = \rho(\pi_1)$. Note again that there are $blk(x_1)$ items before x_1. If $blk(x_1) < a(\pi, 1)$ holds, then B obtains x_1 at her $blk(x_1) + 1$-st action, which violates the assumption that A can obtain X containing x_1 by $\delta(\pi)$.

As the induction step, assuming that the statement holds for $r = t$, we consider the case of $r = t + 1$. If $blk(x_{t+1}) < a(\pi, t + 1)$ holds, then there are $blk(x_{t+1})$ items, as well as $x_1, \ldots, x_t$, before x_{t+1}. From the assumption that $a(\pi, t) \leq blk(x_t)$, A obtains $x_1, \ldots, x_t$ before B's $blk(x_{t+1}) + 1$-st action. Thus, B obtains x_{t+1} by her $blk(x_{t+1}) + 1$-st action, which violates the assumption that $\delta(\pi)$ achieves X. □

Using this result, we can show the following lemma, which is an essential component in the proofs of Lemmas 4 and 5.

LEMMA 3. *Consider any ordering $\rhd$ and any sequence π. Assume that there exists feasible plan δ that achieves bundle X under π and $\rhd$. Then, for any $r' \in \{1, \ldots, m\}$, there also exists feasible plan δ' that achieves X under $\rhd$ and π s.t. for all $r \in \{1, \ldots, m\}$:*

$$\pi'_r = \begin{cases} AB & \text{if } r = r' \\ \pi_r & \text{otherwise.} \end{cases} \tag{1}$$

PROOF. Obviously we can focus on r' such that $r' \leq |X|$. From Corollary 1 and the assumption that there exists $\delta(\pi)$ that achieves X under $\rhd$, a regular induced plan that corresponds to $L_{X,\rhd}$ also achieves X under $\rhd$ and π. From Proposition 3, it holds that for any $r' \in \{1, \ldots, |X|\}$,

$$blk(x_r) \geq r + \rho(\pi_r) - 1.$$

Also, by the definitions of ρ and π', it holds that for any $r \in \{1, \ldots, |X|\}$,

$$r + \rho(\pi_r) - 1 \geq r + \rho(\pi'_r) - 1.$$

By summing them up, we have

$$blk(x_r) \geq r + \rho(\pi'_r) - 1$$

for all $r \in \{1, \ldots, |X|\}$, which coincides with the condition described in Proposition 3. □

LEMMA 4. *Given $\rhd$ and $X \subset \mathcal{O}$, X is possible under $\rhd$ if and only if X can be achieved under $\rhd$ and $\bar{\pi} = sAB \ldots AB$.*

PROOF. The "if" part is obvious. Let us focus on showing the "only if" part. Assume that feasible plan δ achieves X under $\rhd$ for at least one specific sequence π. X can be achieved under $\rhd$ and π. Here, by repeatedly applying a modification described in Eq. 1 to π, we can reach $\bar{\pi}$. Thus, from Corollary 1 and Lemma 3, X can be achieved under $\rhd$ and $\bar{\pi}$. □

LEMMA 5. *Given $\rhd$ and $X \subset \mathcal{O}$, X is certain under $\rhd$ if and only if X can be achieved under $\rhd$ and $\underline{\pi} = sBA \ldots BA$.*

PROOF. The "only if" part is obvious by the definition of certainty of bundle X, Now we turn to show the "if" part. Consider arbitrary bundle X that can be achieved under $\rhd$ and $\underline{\pi}$. From Corollary 1, a regular induced plan $\delta(\underline{\pi})$ achieves X under $\rhd$. From a similar argument with the proof of Lemma 4, X can be achieved by a regular induced plan $\delta'(\pi)$ under $\rhd$, where π is any possible sequence, which guarantees that X is certain under $\rhd$. □

REFERENCES

[1] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, and T. Walsh. Manipulating the probabilistic serial rule. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*, pages 1451–1459, 2015.

[2] H. Aziz, T. Walsh, and L. Xia. Possible and necessary allocations via sequential mechanisms. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15)*, pages 468–474, 2015.

[3] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

[4] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 73–78, 2011.

[5] S. Bouveret and J. Lang. Manipulating picking sequences. In *Proceedings of the Twenty-First European Conference on Artificial Intelligence (ECAI-14)*, pages 141–146, 2014.

[6] S. J. Brams, M. Kilgour, and C. Klamler. Two-person fair division of indivisible items: An efficient, envy-free algorithm. *Notices of the AMS*, 61(2):130–141, 2014.

[7] E. B. Budish and E. Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, 102(5):2237–71, 2012.

[8] E. Fujita, J. Lesca, A. Sonoda, T. Todo, and M. Yokoo. A complexity approach for core-selecting exchange with multiple indivisible goods under lexicographic preferences. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, pages 907–913, 2015.

[9] W. Huang, J. Lou, and Z. Wen. A parallel elicitation-free protocol for allocating indivisible goods. In *Seventh Multidisciplinary Workshop on Advances in Preference Handling (MPREF-13)*, 2013.

[10] T. Kalinowski, N. Narodytska, and T. Walsh. A social welfare optimal sequential allocation procedure. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence (IJCAI-13)*, pages 227–233, 2013.

[11] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI-13)*, pages 452–458, 2013.

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

[13] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Manipulation complexity and gender neutrality in stable marriage procedures. *Autonomous Agents and Multi-Agent Systems*, 22(1):183–199, 2011.

[14] R. Vetschera and D. M. Kilgour. Fair division of indivisible items between two players: design parameters for contested pile methods. *Theory and Decision*, 76(4):547–572, 2014.