

Computing a Small Agreeable Set of Indivisible Items^{*}

Pasin Manurangsi
UC Berkeley

Warut Suksompong
Stanford University

We study the problem of assigning a small subset of indivisible items to a group of agents so that the subset is *agreeable* to all agents, meaning that all agents value the subset as least as much as its complement. For an arbitrary number of agents and items, we derive a tight worst-case bound on the number of items that may need to be included in such a set. We then present polynomial-time algorithms that find an agreeable set whose size matches the worst-case bound when there are two or three agents. We also show that finding small agreeable sets is possible even when we only have access to the agents’ preferences on single items. Furthermore, we investigate the problem of efficiently computing an agreeable set whose size approximates the size of the smallest agreeable set for any given instance. We consider two well-known models for representing the preferences of the agents—the value oracle model and additive utilities—and establish tight bounds on the approximation ratio that can be obtained by algorithms running in polynomial time in each of these models.

1 Introduction

A typical resource allocation problem involves dividing a set of resources among interested agents. We are often concerned with the *efficiency* of the allocation, e.g., achieving high social welfare or ensuring that no other allocation would make every agent better off than the current allocation. Another important issue is the *fairness* of the allocation. For example, we might want the resulting allocation to be *envy-free*, meaning that every agent regards her bundle as the best among all bundles in the allocation [23, 54], or *proportional*, meaning that every agent obtains at least her proportionally fair share [51]. A common feature of such problems is that one agent’s gain is another agent’s loss: The setting inherently puts the agents in conflict with one another, and our task is to try

^{*}This paper unifies and expands earlier versions that appeared in Proceedings of the 25th International Joint Conference on Artificial Intelligence [52] and Proceedings of the 26th International Joint Conference on Artificial Intelligence [40]. In particular, Theorems 5, 10, and 11 are new to this version, and Theorems 4, 6, and 7 improve corresponding results in the earlier versions. These additions lead to asymptotically tight bounds in Sections 3.3 and 4.1.

to resolve this conflict as best we can according to our objectives. Resource allocation problems constitute a major area of study in artificial intelligence.

We study in this work a variant of the resource allocation problem where instead of the agents being pitted against one another, they belong to one and the same group. We will collectively allocate a subset of items to this group; our goal is to make this subset “agreeable” to all agents. Agreeability can be thought of as a minimal desirability condition: While an agent may be able to find other subsets of items that she personally prefers, the current subset is still acceptable for her, and she can agree with its allocation to the group. In other words, if the agreeability condition is not met for some agent, then the agent will be unsatisfied and tempted to leave the group. We consider a notion of agreeability based on the fairness notion of envy-freeness: a subset of items is said to be *agreeable* to an agent if the agent likes it at least as much as the complement set. Agreeability, or minor variants thereof, has been considered in the context of fair division, where each group consists of a single agent [7, 13, 16]. For example, Brams et al. [16] calls the property “worth at least 50 percent”. An appealing aspect of agreeability is that it can be defined for arbitrary ordinal preferences, which constitutes a considerably larger class of preferences than those represented by additive cardinal utility functions [5, 15, 17, 39, 53]. Indeed, for most of this work we only assume that the agents’ preferences are *monotonic*, meaning that an agent always values a set of items at least as much as any of its subsets. Since in a large majority of resource allocation settings agents can simply ignore items that yield negative value to them, the monotonicity assumption is usually made without loss of generality.

As applications of our agreeability notion, one could imagine that the agents are going together on a trip and agreeing on the set of items to put in a shared luggage, or choosing a subset of items as prizes from a team competition that they won together. Without further constraints, the problem described so far would be trivial, since we could simply allocate the entire set of items to the agents. We therefore impose a constraint that the allocated subset should be as small as possible. This constraint on size is reasonable in a variety of settings, including in the two given examples. Indeed, in the first example a luggage has limited space, and in the second example the organizers may want some items to be left as prizes for the losing teams, perhaps so that the allocation seems fair to an outsider. Similar cardinality constraints have been considered in the context of fair division [11]. In the example of agents going together on a trip, a subset of items that they take is agreeable if they like it no less than the complement subset of items left at home. Put differently, based on the set of items chosen, every agent would rather go on the trip than stay at home. Similarly, for agents taking items as prizes from a team competition, if the competition is between two teams and a subset of items is not agreeable to some agent in the winning team, we will have an undesirable situation where the agent envies the losing team that takes the remaining items.

While our study is based on the framework of resource allocation, agreeability is also relevant in other areas of social choice theory and artificial intelligence. In particular, one could think of choosing an agreeable set of items as an election of a committee from a set of candidates, where the committee size is unspecified but perhaps should be minimized. The theory of committee elections provides a number of specific ways to instantiate the

	Full preferences	Ordinal preferences
$n = 2$	$\lfloor \frac{m}{2} \rfloor + 1$ (Thm. 1)	$\lfloor \frac{m}{2} \rfloor + 1$ (Thm. 2)
$n \geq 3$	$\min(\lfloor \frac{m+n}{2} \rfloor, m)$ (Thm. 1)	$\frac{m}{2} + \Theta(\log m)$ for constant n (Thm. 4, 5)

Table 1: Summary of the upper bounds on the size of the smallest agreeable set, presented in Section 3.

notion of agreeability. For example, if one uses approval elections, where every agent either approves or disapproves each candidate and approves a committee if it contains at least one of her approved candidates, an agreeable committee according to our notion corresponds to one where every agent has an approved candidate in the committee. In general, the preferences of the agents for various committees can be quite complex, and several variants of committee elections have been investigated in the literature [6, 50]. We see our work as a starting point that deals with a particularly simple and natural agreeability notion, and our hope is that this work will lay a foundation for studying different notions that may be appropriate for other applications.

1.1 Our Results

In this work, we initiate the study of agreeability in resource allocation. First, in Section 3, we establish upper bounds on the size of the smallest agreeable set, both when the algorithm has access to the agents’ full preferences and when the algorithm only has access to the agents’ preferences on single items. In addition, we present algorithms that compute agreeable sets whose size matches the worst-case bounds under both assumptions. Our results in this section are summarized in Table 1.

In Section 3.1, we derive a tight upper bound on the number of items that may need to be included in an agreeable set, for any number of agents and items. Remarkably, even though agents may have vastly differing and perhaps conflicting preferences, the number of extra items that we might need to choose in order to accommodate all of them is surprisingly small, i.e., half an item per additional agent (Theorem 1). Our result holds under a very weak assumption that preferences are *monotonic*, meaning that an agent cannot be worse off whenever an item is added to her set. Interestingly, to establish this result we make use of Kneser’s conjecture, a combinatorial result whose proof by Lovász [37] gave rise to the field of topological combinatorics.

In Section 3.2, we turn our attention to the question of whether we can efficiently compute an agreeable set whose size matches the worst-case bound given in Section 3.1. We answer the question in the affirmative for the cases of two and three agents. To this end, we make the assumption that preferences are *responsive*, meaning that an agent cannot be worse off when an item is added to her set or replaced by another item that she weakly prefers to the original item. While responsiveness is stronger than monotonicity, it is still a generalization of additivity, a very common assumption on preferences in resource allocation problems. We present polynomial-time algorithms that compute an

agreeable subset whose size matches the upper bound when there are two or three agents (Theorems 2 and 3).

In Section 3.3, we assume that the algorithm only has access to the agents' ordinal preferences on single items rather than subsets of items. Models of this type offer the advantage that the associated algorithms are often simple to implement and the agents do not need to give away or even determine their entire preferences; such models have therefore received widespread attention [7, 13, 31]. With only the ordinal preferences on single items at its disposal, however, in most cases the algorithm cannot tell whether a certain subset is agreeable to an agent or not. Nevertheless, by assuming that preferences are responsive, we can extend preferences on single items to partial preferences on subsets. This allows us to deduce that certain subsets are always agreeable as long as the full responsive preferences are consistent with the rankings over single items; we call such subsets *necessarily agreeable*. Denoting by m the number of items, we show using results from discrepancy theory that for any constant number of agents, there exists a necessarily agreeable subset of size $m/2 + O(\log m)$, and such a subset can be found in polynomial time (Theorem 4). Furthermore, we establish the tightness of this bound by showing that even with three agents, there exist preferences for which every necessarily agreeable subset has size $m/2 + \Omega(\log m)$ (Theorem 5).

Next, in Section 4, we investigate the problem of computing an agreeable subset of approximately optimal size for any given instance, as opposed to one whose size matches the worst-case bound over all instances with the same number of agents and items. We tackle the problem using two models for representing preferences that are well-studied in the literature, and exhibit computationally efficient algorithms for finding an agreeable set of approximately optimal size in each of them. Moreover, in both of the models we show that our approximation ratios are asymptotically tight.

In Section 4.1, we consider general preferences using the value oracle model, where the preferences of the agents are represented by utility functions and the algorithm is allowed to query the utility of any agent for any subset of items. We exhibit an efficient approximation algorithm with approximation ratio $O(m/\log m)$ in this model (Theorem 7). While this may not seem impressive, especially in light of the observation that the trivial algorithm which always outputs the entire set of items already achieves approximation ratio $O(m)$, we show that our approximation ratio is in fact the best we can hope for. In other words, there does not exist a polynomial time algorithm with approximation ratio $o(m/\log m)$, even when there is only a single agent (Theorem 8).

In Section 4.2, we assume that the agents' preferences are represented by additive utility functions. Additivity provides a reasonable tradeoff between simplicity and expressiveness; it is commonly assumed in the literature, especially in recent work [5, 15, 17, 39, 53]. We show that under additive valuations, it is NP-hard to decide whether there exists an agreeable set containing exactly half of the items, even where there are only two agents (Theorem 9). On the other hand, using results on covering integer programs, we demonstrate the existence of an $O(\log n)$ -approximation algorithm for computing a minimum size agreeable set (Theorem 13). Moreover, we show that this approximation factor is tight: For any constant $\delta > 0$, it is NP-hard to approximate the problem to within a factor of $(1 - \delta) \ln n$ (Theorem 12).

1.2 Related Work

While resource allocation and fair division have been extensively studied in the artificial intelligence literature [7, 11, 13, 14, 15, 18, 29, 33], the group aspect that we consider in this work is relatively new. Skowron et al. [50] studied a similar setting where a group of agents collectively decide on a common set of items. In their work, the number of items that can be chosen is fixed and the agents strive to maximize their utility subject to that constraint, whereas in our work the number of items is variable but should be minimized, and an agreeability constraint is placed on the chosen set. A related problem that has recently received attention is that of fairly allocating items to groups of agents; the problem has been studied both in the context of indivisible items [39, 48, 53] and divisible items [46, 47]. Like in our work, the items there are treated as public goods within each group—all members of a group derive full utility from the items allocated to the group. Unlike our work, however, all of the works mentioned above assume that there are multiple groups and the entire resource should be allocated to the groups. After the publication of the earlier versions of our work [40, 52], Gourvès [25] studied agreeable sets under the additional assumption that the sets must satisfy matroidal constraints.

In addition to the literature on resource allocation and fair division, another line of work relevant to ours is that of combinatorial voting [1, 34, 55]. A typical example of combinatorial voting is an election in which voters are asked to decide on a certain set of issues, and their preferences on the outcomes of various issues are dependent. For instance, a voter might be in favor of each proposed policy individually, but thinks that too much money would be spent if all of the policies were to be implemented. Combinatorial voting can be framed in our resource allocation setting, with the items to be allocated corresponding to the issues being voted upon. Several aspects of combinatorial voting has been studied, including the communication and computational cost of different voting rules, implementation methods, and strategic behavior of the voters.

2 Preliminaries

We consider n agents, numbered $1, 2, \dots, n$, who will be collectively allocated a subset of the set $S = \{x_1, x_2, \dots, x_m\}$ of m indivisible items. Denote by \mathcal{S} the set of all subsets of S . Each agent i is endowed with a preference relation \succeq_i , a reflexive, complete, and transitive ordering over \mathcal{S} . Let \succ_i denote the strict part and \sim_i the indifference part of the relation \succeq_i . For items x and y , we will sometimes abuse notation and write $x \succeq y$ to mean $\{x\} \succeq \{y\}$.

We assume in this paper that preferences are monotonic, i.e., an agent cannot be worse off when an item is added to her set. Monotonicity is a natural assumption in a wide range of situations. In particular, it implies free disposal of items—every item is considered to be of nonnegative value to each agent.

Definition 1. *A preference \succeq on \mathcal{S} is monotonic if $T \cup \{x\} \succeq T$ for all $T \subseteq S$.*

Note that if $x \in T$, then $T \cup \{x\} \succeq T$ always holds, so we only need to check the condition when $x \in S \setminus T$.

We are now ready to define the central notion of this paper.

Definition 2. A subset $T \subseteq S$ is said to be agreeable to agent i if $T \succeq_i S \setminus T$.

When the set of agents considered is clear from the context, we will sometimes refer to a set that is agreeable to all agents simply as an agreeable set. Since preferences are monotonic, the whole set S is agreeable to every agent, so an agreeable set always exists for any number of agents.¹ Agreeability to an agent also implies that the agent does not strictly prefer any subset of the complement to the current set. That is, we have $T \succeq_i U$ for any $U \subseteq S \setminus T$.

Another property of preferences that we will consider is responsiveness, which says that an agent cannot be worse off whenever an item is added to her set or replaced by another item that she weakly prefers to the original item. While stronger than monotonicity, responsiveness is still a reasonable assumption in many settings.²

Definition 3. A preference \succeq on \mathcal{S} is responsive if it satisfies the following two conditions:

- \succeq is monotonic;
- $(T \setminus \{y\}) \cup \{x\} \succeq T$ for all $T \subseteq S$ and $x, y \in S$ such that $x \succeq y$, $x \notin T$ and $y \in T$.

If we have access to the complete preference of an agent, we can check whether a subset is agreeable to the agent simply by comparing it to its complement. When we only have access to the agent's preference on single items, however, most of the time we cannot tell whether a given subset is agreeable or not. Nevertheless, if we assume that the agent's preference is responsive, we can sometimes deduce that a certain subset is agreeable only by looking at the agent's preference on single items. The following definition captures this intuition. In general, we use \succeq to denote a preference on \mathcal{S} and \succeq^{sing} to denote a preference on the single items in S .

Definition 4. Fix a preference \succeq^{sing} on the single items in S . A subset $T \subseteq S$ is said to be necessarily agreeable with respect to \succeq^{sing} if $T \succeq S \setminus T$ for any responsive preference \succeq on \mathcal{S} consistent with \succeq^{sing} .

For the sake of brevity, we say that a subset of items is necessarily agreeable to an agent if it is necessarily agreeable with respect to the preference on single items of the agent.

We now make a connection to the model where every agent has a cardinal utility for each subset of items. A *utility function* u is a function that maps any subset of items to a nonnegative real number. Since each agent's preference is reflexive, complete, and transitive, there is a utility function $u_i : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ for agent i such that for any $T_1, T_2 \subseteq S$, we have $T_1 \succeq_i T_2$ if and only if $u_i(T_1) \geq u_i(T_2)$. Moreover, since we consider monotonic

¹If preferences are not monotonic, an agreeable set might not exist, e.g., if there are two agents with strict preferences, and one agent's preference is exactly the opposite of the other agent's preference.

²For a comprehensive treatment of properties concerning the ranking of sets of objects, we refer to a survey by Barberà et al. [9].

preferences, we have $u_i(T_1) \leq u_i(T_2)$ for any $T_1 \subseteq T_2$. We assume that $u_i(\emptyset) = 0$ for all i . A utility function u is said to be *additive* if $u(T_1 \cup T_2) = u(T_1) + u(T_2)$ for any disjoint subsets T_1, T_2 , and *subadditive* if $u(T_1 \cup T_2) \leq u(T_1) + u(T_2)$ for any T_1, T_2 . Any monotonic additive function is also subadditive. Subadditive utility functions have been extensively studied in the literature [10, 21].

When the preferences of the agents are given by subadditive utility functions, a subset that an agent regards as agreeable also gives the agent a utility of at least half of the agent's utility for the whole set S . Indeed, for any agreeable subset T we have

$$f(S) = f(T \cup (S \setminus T)) \leq f(T) + f(S \setminus T) \leq 2f(T),$$

which implies that $f(T) \geq f(S)/2$. Hence an agreeable subset also gives a 2-approximation of the welfare to an agent when the agent's utility function is subadditive.

We end this section by giving a characterization of necessarily agreeable subsets, which will be used multiple times in the paper. Similar statements have been shown by Aziz et al. [7] and Brams et al. [16], although our treatment differs slightly in dealing with ties.

Proposition 1. *Fix a preference \succeq^{sing} on the single items in S with*

$$x_1 \succeq^{sing} x_2 \succeq^{sing} \dots \succeq^{sing} x_m.$$

Let $T \subseteq S$, and define $I_k = \{x_1, x_2, \dots, x_k\}$ for all $k = 1, 2, \dots, m$.

If $|I_k \cap T| \geq k/2$ for all $k = 1, 2, \dots, m$, then T is necessarily agreeable with respect to \succeq^{sing} . The converse also holds if the preference \succeq^{sing} is strict.

Proof. Assume first that $|I_k \cap T| \geq k/2$ for all $k = 1, 2, \dots, m$. Since $|I_m \cap T| \geq m/2$, we have that $|T| \geq |S \setminus T|$. Let $T' \subseteq T$ be the subset consisting of the $|S \setminus T|$ items of T with the smallest indices.

Define a bijective function $f : T' \rightarrow S \setminus T$ as follows: Given the item $x_k \in T$ with the smallest index for which $f(x_k)$ is not yet defined, we define $f(x_k)$ to be the item in $S \setminus T$ with the smallest index that has not occurred in the range of f so far. Since $|I_k \cap T| \geq k/2$ for all $k = 1, 2, \dots, m$, the function f maps each item x_k to another item x_l with $l > k$. The definition of responsiveness implies that for any responsive preference \succeq on \mathcal{S} consistent with \succeq^{sing} , it holds that $T' \succeq S \setminus T$. Since any responsive preference is also monotonic, we have $T \succeq S \setminus T$, which implies that T is necessarily agreeable with respect to \succeq^{sing} .

For the converse, assume that the preference \succeq^{sing} is strict, and that $|I_l \cap T| < l/2$ for some $l = 1, 2, \dots, m$. Let $\epsilon > 0$ be a small constant, and suppose that the preference \succeq is given by an additive utility function u such that:

- $u(x_i) = 1 + (l - i)\epsilon$ for $1 \leq i \leq l$;
- $u(x_i) = (m - i)\epsilon$ for $l < i \leq m$.

Since any preference that can be represented by an additive utility function is responsive, \succeq is responsive. Moreover, we have $u(S \setminus T) > l/2$, whereas $u(T) < l/2$ when ϵ is small enough. It follows that \succeq is a responsive preference on \mathcal{S} consistent with \succeq^{sing} such that $S \setminus T \succ T$. Hence T is not necessarily agreeable with respect to \succeq^{sing} . \square

Finally, any logarithm written without a base in this paper is assumed to have base 2.

3 Worst-Case Bounds

In this section, we establish upper bounds on the size of the smallest agreeable set, both when the algorithm has access to the agents' full preferences and when the algorithm only has access to the agents' preferences on single items. In addition, we present algorithms that compute agreeable sets whose size matches the worst-case bounds under both assumptions.

3.1 General Worst-Case Bound

We commence our study of agreeable sets by deriving a tight worst-case bound on the number of items that may need to be included in such a set, for any number of items and any number of agents with arbitrary preferences on the items. Even with a single agent, there already exists a preference for which we need to include at least half of the items, e.g., a preference represented by an additive utility function that gives the same positive utility to every item. In light of this, it may seem that there is little hope of obtaining a small agreeable set when there are several agents, possibly with wildly varying preferences. Nevertheless, we show that the number of extra items that we need to include to accommodate the additional agents is surprisingly small even in the worst case—this number is only half an item per additional agent.

Theorem 1. *For any number of agents and items, there exists a subset $T \subseteq S$ such that*

$$|T| \leq \min \left(\left\lfloor \frac{m+n}{2} \right\rfloor, m \right)$$

and T is agreeable to all agents. Moreover, there exist preferences for which the bound $\min \left(\left\lfloor \frac{m+n}{2} \right\rfloor, m \right)$ is tight.

Theorem 1 can be seen as a discrete version of consensus halving, where the goal is to partition a divisible item such as cake or land into two parts that all agents think are worth exactly the same. Interestingly, a consensus halving partition can be found for any number of agents [2, 49]. It follows that we can find a part of the item that is at most half of the item but that all agents think is worth at least half of the item (in particular, we choose the smaller of the two parts in the consensus halving partition). When items are indivisible, however, it may no longer be possible to choose a set containing at most half of the items such that all agents believe this set is worth at least as much as its complement. Indeed, if there is only one item and all agents value this item positively, the item needs to be included in the set. Theorem 1 gives us a precise bound on how many additional items need to be included in the worst case.

We first give a direct proof of Theorem 1 for the case of two agents; our proof for the general case will rely on a combinatorial result.

Direct proof of Theorem 1 when $n = 2$. Denote by \succeq_1 and \succeq_2 the preferences on S of the two agents. We establish the existence of a set of size at most $\lfloor \frac{m+2}{2} \rfloor$ that is agreeable to both agents; the tightness of the bound follows in the same way as in our proof of Theorem 1 for any number of agents.

Assume first that $m = 2k + 1$ is odd. Suppose for contradiction that no subset of size at most $k + 1$ is agreeable to both agents. Let $T \subseteq S$ be such that $|T| = k$. We begin by proving the following claim.

Claim: If $T \succ_1 S \setminus T$, then

$$(T \cup \{x\}) \setminus \{x'\} \succ_1 ((S \setminus T) \setminus \{x\}) \cup \{x'\}$$

for any $x \in S \setminus T$ and $x' \in T$.

Proof of Claim: Suppose that $T \succ_1 S \setminus T$, $x \in S \setminus T$, and $x' \in T$. It follows from monotonicity that $T \cup \{x\} \succ_1 (S \setminus T) \setminus \{x\}$. Since no subset of size $k + 1$ is agreeable to both agents, we have $(S \setminus T) \setminus \{x\} \succ_2 T \cup \{x\}$. By monotonicity again, we have

$$((S \setminus T) \setminus \{x\}) \cup \{x'\} \succ_2 (T \cup \{x\}) \setminus \{x'\}.$$

Using again the assumption that no subset of size $k + 1$ is agreeable to both agents, it follows that

$$(T \cup \{x\}) \setminus \{x'\} \succ_1 ((S \setminus T) \setminus \{x\}) \cup \{x'\},$$

and our claim is proved. \square

We now use our claim to obtain the desired contradiction. Assume without loss of generality that $\{x_1, x_2, \dots, x_k\} \succ_1 \{x_{k+1}, x_{k+2}, \dots, x_{2k+1}\}$. Applying our claim repeatedly to move items between the two sets, we find that

$$\{x_{k+1}, x_2, \dots, x_k\} \succ_1 \{x_1, x_{k+2}, \dots, x_{2k+1}\},$$

$$\{x_{k+1}, x_{k+2}, x_3, \dots, x_k\} \succ_1 \{x_1, x_2, x_{k+3}, \dots, x_{2k+1}\},$$

and so on, until finally

$$\{x_{k+1}, x_{k+2}, \dots, x_{2k}\} \succ_1 \{x_1, x_2, \dots, x_k, x_{2k+1}\}.$$

By monotonicity, we have $\{x_{k+1}, x_{k+2}, \dots, x_{2k+1}\} \succ_1 \{x_1, x_2, \dots, x_k\}$, which contradicts our assumption that $\{x_1, x_2, \dots, x_k\} \succ_1 \{x_{k+1}, x_{k+2}, \dots, x_{2k+1}\}$.

Assume now that $m = 2k$ is even. Let S' be the set of all items in S except x_1 . We know from the case where m is odd that there exists a subset $T \subseteq S'$ of size at most k such that $T \succeq_1 S' \setminus T$ and $T \succeq_2 S' \setminus T$. Since preferences are monotonic, we have that $T \cup \{x_1\} \succeq_1 S' \setminus T$ and $T \cup \{x_1\} \succeq_2 S' \setminus T$. This means that the set $T \cup \{x_1\}$ of size at most $k + 1$ is our desired subset, completing the proof.

Note that this proof also yields a polynomial-time algorithm to compute an agreeable set of size at most $\lfloor \frac{m+2}{2} \rfloor$ that is agreeable to both agents. Assume that $m = 2k + 1$ is odd; the case where m is even can be handled similarly. Let $T \subseteq S$ be an arbitrary subset of size k . If $S \setminus T \succeq_1 T$ and $S \setminus T \succeq_2 T$, we are done. Otherwise, assume without loss of generality that $T \succ_1 S \setminus T$, and choose arbitrarily $x \in S \setminus T$ and $x' \in T$. As

in the proof of the claim, if $T \cup \{x\} \succeq_2 (S \setminus T) \setminus \{x\}$, or if $(S \setminus T) \setminus \{x\} \succ_2 T \cup \{x\}$ and $((S \setminus T) \setminus \{x\}) \cup \{x'\} \succeq_1 (T \cup \{x\}) \setminus \{x'\}$, we are done. Hence we may assume as in the conclusion of the claim that $(T \cup \{x\}) \setminus \{x'\} \succ_1 ((S \setminus T) \setminus \{x\}) \cup \{x'\}$. This means that we can find an agreeable subset by moving elements repeatedly between the two sets as in the continuation of the proof. Since we need to move elements at most k times, our algorithm runs in polynomial time. \square

We now move on to the general case. As alluded earlier, our proof of the theorem relies on the following combinatorial result, which is best known as Kneser's conjecture.

Lemma 1 (Kneser's conjecture). *Let G be the undirected graph with all k -element subsets of the set $\{1, 2, \dots, n\}$ as vertices such that there exists an edge between two vertices if and only if the corresponding sets are disjoint. The chromatic number³ of G is given by*

$$\chi(G) = \begin{cases} n - 2k + 2 & \text{if } n \geq 2k; \\ 1 & \text{otherwise.} \end{cases}$$

The statement of the lemma is due to Kneser [30], who proposed it as a conjecture in the problem column of a German mathematics magazine in 1955. In spite of the simple statement, it was not until 1978 that the conjecture was first resolved by Lovász [37] using topological methods. The proof was later simplified by Bárány [8] and Greene [26], before Matoušek [41] gave the first purely combinatorial proof in 2004. Lovász's proof of the conjecture, which makes use of the Borsuk-Ulam theorem, marked the first time that methods from algebraic topology were used to solve a combinatorial problem, and gave rise to the field of topological combinatorics.

With Lemma 1 in hand, we are ready to establish our theorem.

Proof of Theorem 1. Let $k = \lfloor \frac{m+n}{2} \rfloor$. If $k \geq m$, the set S of all items has size $m = \min(k, m)$ and is agreeable to all agents since preferences are monotonic. Assume from now on that $k < m$, and consider the undirected graph G with all $(m-k)$ -element subsets of $\{x_1, x_2, \dots, x_m\}$ as vertices and with edges connecting vertices whose corresponding sets are disjoint. For example, if $n = 2$ and $m = 5$ (so $k = 3$ and $m - k = 2$), the graph G corresponds to the well-known Petersen graph and is shown in Figure 1.

If all agents weakly prefer $S \setminus T$ to T for some $(m-k)$ -element subset $T \subseteq S$, then $S \setminus T$ is our desired agreeable subset of size k . Suppose for contradiction that for any $(m-k)$ -element subset $T \subseteq S$, there exists an agent who strictly prefers T to $S \setminus T$. We color the vertices of G with n colors in the following way: For each vertex v of G corresponding to a set T , color it with the color corresponding to an agent who strictly prefers T to $S \setminus T$. If there is more than one such agent, choose one arbitrarily.

Since $k = \lfloor \frac{m+n}{2} \rfloor \geq \frac{m}{2}$, we have $m \geq 2(m-k)$. By Lemma 1, the chromatic number of G is

$$m - 2(m-k) + 2 = 2k - m + 2 \geq 2 \left(\frac{m+n-1}{2} \right) - m + 2 = n + 1.$$

³The chromatic number of a graph is defined as the smallest number of colors needed to color the vertices of the graph so that no two adjacent vertices share the same color.




Figure 1: The graph G in the proof of Theorem 1 when $n = 2$ and $m = 5$, also known as the Petersen graph. A vertex with a label $\{i, j\}$ corresponds to the set $\{x_i, x_j\}$.

Since we colored G with n colors, there exist two adjacent vertices sharing the same color. Let T_1 and T_2 be the sets corresponding to these vertices. This means that $T_1 \succ_i S \setminus T_1$ and $T_2 \succ_i S \setminus T_2$ for some agent i . Since T_1 and T_2 are disjoint, we have $T_1 \subseteq S \setminus T_2$ and $T_2 \subseteq S \setminus T_1$. Monotonicity now implies that

$$S \setminus T_1 \succeq_i T_2 \succ_i S \setminus T_2 \succeq_i T_1 \succ_i S \setminus T_1,$$

a contradiction. This means that we can always find a subset of size k that is agreeable to all agents.

Finally, we show that there exist monotonic preferences for which the bound $\min(k, m)$ is tight. In fact, we can even choose preferences that are represented by additive utility functions. We consider two cases.

- $n \geq m$. Then $\min(k, m) = m$. For $i = 1, 2, \dots, n$, let the preference of agent i be given by an additive utility function u such that $u(x_{\min(i, m)}) = 1$ and $u(x_j) = 0$ for all $j \neq \min(i, m)$. Then any subset that is agreeable to agent i must contain item $x_{\min(i, m)}$. Hence a subset that is agreeable to all agents must contain all m items.
- $n < m$. Then $\min(k, m) = k$. For $i = 1, 2, \dots, n - 1$, let the preference of agent i be given by an additive utility function u such that $u(x_i) = 1$ and $u(x_j) = 0$ for all $j \neq i$. Let the preference of agent n be given by an additive utility function u such that $u(x_j) = 1$ for $j \in \{n, n + 1, \dots, m\}$ and $u(x_j) = 0$ otherwise.

For $i = 1, 2, \dots, n - 1$, any subset that is agreeable to agent i must contain item x_i . Also, any subset that is agreeable to agent n must contain at least half of the items x_n, x_{n+1}, \dots, x_m . Hence a subset that is agreeable to all agents must have size at least $n - 1 + \lceil \frac{m-n+1}{2} \rceil = \lceil \frac{m+n-1}{2} \rceil = \lfloor \frac{m+n}{2} \rfloor = k$, as desired.

This completes the proof. \square

3.2 Matching the Worst-Case Bound

Theorem 1 gives us a tight worst-case bound on the size of the smallest agreeable set for any number of agents and items. However, its proof does not yield a method for obtaining a set of that size. Since the number of sets that we have to consider is exponential in the number of items, brute-force search is infeasible even for moderate numbers of items. Our goal in this section is to show that when there are two or three agents with responsive preferences, it is in fact possible to compute an agreeable set whose size matches the worst-case bound in polynomial time. This implies that we can compute such a set even when the number of items is large.

An important issue when we discuss algorithms is how we represent the agents' preferences. Since preferences on subsets, unlike preferences on single items, might not have a succinct representation, it is not possible to design algorithms that run in time polynomial in the number of items if the algorithm is required to read the entire preference. To circumvent this problem, we assume in this section that preferences are responsive; this allows us to extend preferences on single items to partial preferences on subsets. Our algorithm for two agents will only make use of the preferences on single items and compute a necessarily agreeable subset.⁴ On the other hand, our algorithm for three agents will also query the agents' preferences on subsets through a preference oracle in addition to utilizing the preferences on single items.

We first handle the case of two agents.

Theorem 2. *Assume that there are two agents with preferences \succeq_1^{sing} and \succeq_2^{sing} on the single items in S . There exists a subset $T \subseteq S$ such that $|T| \leq \lfloor \frac{m+2}{2} \rfloor$ and T is necessarily agreeable with respect to both \succeq_1^{sing} and \succeq_2^{sing} . Also, there exists a polynomial-time algorithm that computes such a subset T .*

Moreover, there exist preferences on the single items in S for which the bound $\lfloor \frac{m+2}{2} \rfloor$ is tight.

Proof. Assume first that $m = 2k + 1$ is odd, and suppose without loss of generality that $x_1 \succeq_1^{sing} x_2 \succeq_1^{sing} \dots \succeq_1^{sing} x_{2k+1}$. We choose our set T of $\lfloor \frac{m+2}{2} \rfloor = k + 1$ items as follows:

1. Choose x_1 .
2. Between each of the k pairs of items $(x_2, x_3), (x_4, x_5), \dots, (x_{2k}, x_{2k+1})$, choose the item that is preferred according to \succeq_2^{sing} . If \succeq_2^{sing} is indifferent between any pair of items, choose an arbitrary item from that pair.

For any $j = 1, 2, \dots, m$, our set T contains at least $j/2$ of the j items x_1, x_2, \dots, x_j ; by Proposition 1, T is necessarily agreeable with respect to \succeq_1^{sing} . Moreover, since

⁴If we do not assume responsiveness, there still exists a polynomial-time algorithm for two agents that discovers the agents' preferences on subsets through a preference oracle; this algorithm is described in the direct proof of Theorem 1 for two agents.

we choose the item that is preferred according to \succeq_2^{sing} from each of the sets $\{x_2, x_3\}$, $\{x_4, x_5\}, \dots, \{x_{2k}, x_{2k+1}\}$ along with x_1 , Proposition 1 implies that T is also necessarily agreeable with respect to \succeq_2^{sing} . Hence T is necessarily agreeable with respect to both \succeq_1^{sing} and \succeq_2^{sing} .

Assume now that $m = 2k$ is even. Let $S' = S \setminus \{x_1\}$. We apply the algorithm from the case where m is odd to choose a set $T \subseteq S'$ of size k that is necessarily agreeable with respect to both \succeq_1^{sing} and \succeq_2^{sing} when the universe considered is S' . It follows that $T \cup \{x_1\}$ is a subset of size $\lfloor \frac{m+2}{2} \rfloor = k + 1$ that is necessarily agreeable with respect to both \succeq_1^{sing} and \succeq_2^{sing} when the universe considered is S .

Next, we show that there exist preferences on single items for which the bound $\lfloor \frac{m+2}{2} \rfloor$ is tight. If $m = 2k + 1$ is odd and the preference \succeq_1^{sing} is strict, then by Proposition 1, any subset that is necessarily agreeable with respect to \succeq_1^{sing} alone must already contain at least $\lfloor \frac{m+2}{2} \rfloor = k + 1$ items.

Finally, suppose that $m = 2k$ is even, and let \succeq_1^{sing} and \succeq_2^{sing} be such that $x_1 \succ_1^{sing} x_2 \succ_1^{sing} \dots \succ_1^{sing} x_{2k}$ and $x_{2k} \succ_2^{sing} x_{2k-1} \succ_2^{sing} \dots \succ_2^{sing} x_1$. By Proposition 1, any subset $T \subseteq S$ that is necessarily agreeable with respect to \succeq_1^{sing} alone must contain at least k items, one of which is x_1 . If T contains exactly k items, then it contains exactly $k - 1$ items among x_2, x_3, \dots, x_{2k} . Proposition 1 implies that such a set T is not necessarily agreeable with respect to \succeq_2^{sing} . Hence any subset $T \subseteq S$ that is necessarily agreeable with respect to both \succeq_1^{sing} and \succeq_2^{sing} must contain at least $\lfloor \frac{m+2}{2} \rfloor = k + 1$ items, as desired. \square

At a high level, the algorithm in Theorem 2 bears a resemblance to the ‘‘Trump rule’’, which was proposed by Pruhs and Woeginger [44] for fair division of indivisible items between two agents. Like our algorithm, the Trump rule takes as input the preferences on single items of the two agents. Using our terminology, the rule is guaranteed to produce an allocation with the property that each agent views her bundle as necessarily agreeable, whenever such an allocation exists. The difference between the Trump rule and our algorithm is that the Trump rule produces a partition of the items into two subsets with each agent taking one subset, whereas our algorithm produces a single subset that both agents share.

Observe that in the case of two agents, the upper bound for the size of the smallest necessarily agreeable set (Theorem 2) coincides with the bound for the size of the smallest agreeable set (Theorem 1). This is somewhat surprising because the definition of a necessarily agreeable set only involves preferences on single items, and yet the worst-case bound remains unchanged even if we have access to the full preferences. The following example shows that the same statement ceases to hold when there are three agents.

Example 1. Let $m = 6$, and assume that the preferences on single items of the three agents are as follows:

1. $x_1 \succ_1^{sing} x_4 \succ_1^{sing} x_5 \succ_1^{sing} x_6 \succ_1^{sing} x_2 \succ_1^{sing} x_3$;
2. $x_2 \succ_2^{sing} x_5 \succ_2^{sing} x_6 \succ_2^{sing} x_4 \succ_2^{sing} x_3 \succ_2^{sing} x_1$;

$$3. x_3 \succ_3^{sing} x_6 \succ_3^{sing} x_4 \succ_3^{sing} x_5 \succ_3^{sing} x_1 \succ_3^{sing} x_2.$$

In Example 1, any subset that is necessarily agreeable to all three agents must contain x_1, x_2, x_3 , since each of them is ranked first by some agent. Moreover, choosing only one of x_4, x_5, x_6 does not yield a necessarily agreeable set for the agent who ranks that item fourth. Hence a necessarily agreeable set must contain at least five items. On the other hand, if we have access to the agents' full preferences, Theorem 1 implies that we can find a set of size $\lfloor \frac{6+3}{2} \rfloor = 4$ that is agreeable to all agents.

Therefore, to compute an agreeable set whose size matches the worst-case bound when there are three agents, it is not sufficient to consider preferences on single items only. Nevertheless, if the algorithm has access to the agents' full preferences, it is possible to find such a subset in polynomial time. To access the preferences, the algorithm is allowed to make a polynomial number of queries to a *preference oracle*. In each query, the algorithm can specify an agent and two subsets of items to the preference oracle, and the oracle reveals the preference of that agent between the two subsets.

Theorem 3. *Assume that there are three agents with responsive preferences \succeq_1, \succeq_2 , and \succeq_3 on \mathcal{S} . There exists a polynomial-time algorithm that computes a subset $T \subseteq S$ such that $|T| \leq \lfloor \frac{m+3}{2} \rfloor$ and T is agreeable to all three agents.*

Proof. Assume first that $m = 2k$ is even. Our goal is to find a subset of size $\lfloor \frac{m+3}{2} \rfloor = k+1$ that is agreeable to all three agents. Suppose without loss of generality that x_{2k-1} is the most preferred item according to \succeq_1 , x_{2k} is the most preferred item other than x_{2k-1} according to \succeq_2 , and among the remaining $2k - 2$ items, the preference \succeq_1 ranks them as $x_1 \succeq_1 x_2 \succeq_1 \dots \succeq_1 x_{2k-2}$.

Let $A = \{x_1, x_2, \dots, x_{2k-2}\}$, and consider the pairs $(x_1, x_2), (x_3, x_4), \dots, (x_{2k-3}, x_{2k-2})$. Let B be a set of $k - 1$ items containing an item from each pair that is *not* preferred to the other item in the pair according to \succeq_2 . If \succeq_2 is indifferent between any pair of items, we choose arbitrarily. Responsiveness implies that $A \setminus B \succeq_2 B$.

As long as $A \setminus B \succeq_2 B$, we remove an element from B that was also originally in B , and insert the other item in its pair into B . We must eventually reach a point where $B \succeq_2 A \setminus B$, at the latest after $k - 1$ moves. We consider two cases.

- We have not performed any move. By definition of B , we have that $B \succeq_2 A \setminus B$ and $A \setminus B \succeq_2 B$, and therefore $A \setminus B \sim_2 B$. Since preferences are monotonic, it follows that $(A \setminus B) \cup \{x_{2k}\} \succeq_2 B$ and $B \cup \{x_{2k}\} \succeq_2 A \setminus B$.
- We have performed at least one move. Suppose without loss of generality that in our last move, we inserted x_{2i-1} into and removed x_{2i} from B . Let $C = (A \setminus (B \cup \{x_{2i}\})) \cup \{x_{2i-1}\}$ and $D = (B \setminus \{x_{2i-1}\}) \cup \{x_{2i}\}$, i.e., C and D are the sets $A \setminus B$ and B before the last move, respectively. We have that $C \succ_2 D$ and $B \succeq_2 A \setminus B$, and it follows from monotonicity that $C \cup \{x_{2k}\} \succeq_2 D$ and $B \cup \{x_{2k}\} \succeq_2 A \setminus B$. We claim that at least one of $D \cup \{x_{2k}\} \succeq_2 C$ and $(A \setminus B) \cup \{x_{2k}\} \succeq_2 B$ holds.

Assume for contradiction that $C \succ_2 D \cup \{x_{2k}\}$ and $B \succ_2 (A \setminus B) \cup \{x_{2k}\}$. Responsiveness implies that

$$C \succ_2 D \cup \{x_{2k}\} \succeq_2 B \succ_2 (A \setminus B) \cup \{x_{2k}\} \succeq_2 C,$$

a contradiction. Hence at least one of $D \cup \{x_{2k}\} \succeq_2 C$ and $(A \setminus B) \cup \{x_{2k}\} \succeq_2 B$ holds, as claimed.

In both cases, we can find in polynomial time a subset $E \subseteq A$ of size $k - 1$ containing an item from each of the pairs $(x_1, x_2), (x_3, x_4), \dots, (x_{2k-3}, x_{2k-2})$ such that $E \cup \{x_{2k}\} \succeq_2 A \setminus E$ and $(A \setminus E) \cup \{x_{2k}\} \succeq_2 E$.

We now choose our agreeable set of size $k + 1$ as follows:

1. Choose both x_{2k-1} and x_{2k} .
2. Choose either E or $A \setminus E$ according to which set agent 3 prefers. (If agent 3 is indifferent between the two sets, choose one of them arbitrarily.)

We claim that our chosen set T is agreeable to all three agents. We prove the claim separately for each of the agents.

- For any $j = 1, 2, \dots, m$, the set T contains at least $j/2$ of the j most preferred items according to \succeq_1 . Since \succeq_1 is responsive, Proposition 1 implies that T is necessarily agreeable to agent 1.
- Since $E \cup \{x_{2k}\} \succeq_2 A \setminus E$ and $A \setminus E \cup \{x_{2k}\} \succeq_2 E$, and T contains either E or $A \setminus E$ along with both x_{2k-1} and x_{2k} , T is agreeable to agent 2.
- Since we choose the set E or $A \setminus E$ that agent 3 prefers and we include both of the remaining items x_{2k-1} and x_{2k} , T is agreeable to agent 3.

Hence T is agreeable to all three agents, as claimed. This concludes the analysis of the case where m is even.

Finally, assume that $m = 2k + 1$ is odd. Our goal is to find a subset of size $\lfloor \frac{m+3}{2} \rfloor = k + 2$ that is agreeable to all three agents. Let $S' = S \setminus \{x_1\}$. We apply the algorithm from the case where m is even to choose a set $T \subseteq S'$ of size $k + 1$ that is agreeable to all three agents when the universe considered is S' . It follows that $T \cup \{x_1\}$ is a subset of size $k + 2$ that is agreeable to all three agents when the universe considered is S . \square

3.3 Computing Small Necessarily Agreeable Sets

In this section, we consider a model in which the algorithm only has access to each agent's ranking over the items. We will therefore be interested in computing a small subset that is *necessarily* agreeable to every agent. While the algorithm has significantly less information at its disposal than before, as we will see, it is still possible to find small subsets that are necessarily agreeable to all agents.

If the algorithm had access to the agents' preferences over all subsets of items, Theorem 1 implies that it could always find a subset of size $\lfloor \frac{m+n}{2} \rfloor$ that is agreeable to all agents. For two agents, the algorithm in Theorem 2 only uses the agents' rankings to compute a subset of this size that is necessarily agreeable to both agents. As Example 1 shows, however, a necessarily agreeable subset of this size might not exist even when there are three agents. Indeed, it is not clear how much extra "penalty" we have to pay

for the information restriction that we are imposing. For example, it could be that with three agents, there already exist preferences on single items for which any necessarily agreeable subset contains at least cm items for some constant $c > 1/2$. We show in the next theorem that this is in fact not the case—there always exists a necessarily agreeable subset of size $m/2 + O(\log m)$ as long as the number of agents is constant. Moreover, such a set can be computed in polynomial time.

Theorem 4. *For any constant number of agents, there exists a subset of S of size $m/2 + O(\log m)$ that is necessarily agreeable to all agents. Moreover, such a subset can be computed in polynomial time.*

To prove this theorem, we will use the following result from discrepancy theory due to Bohus [12].

Lemma 2 ([12]). *Let $\sigma_1, \sigma_2, \dots, \sigma_n$ be permutations of the set $M = \{1, 2, \dots, m\}$. There exists a function $f : M \rightarrow \{-1, 1\}$ such that for any $1 \leq p \leq q \leq m$ and any $1 \leq j \leq n$,*

$$\left| \sum_{i=p}^q f(\sigma_j(i)) \right| \leq 8n \log m.$$

Moreover, such a function can be computed in polynomial time.

Proof of Theorem 4. Suppose that agent j ranks the single items as $x_{\sigma_j(1)} \succeq_j^{sing} x_{\sigma_j(2)} \succeq_j^{sing} \dots \succeq_j^{sing} x_{\sigma_j(m)}$. By Lemma 2, we can efficiently compute a function $f : S \rightarrow \{-1, 1\}$ such that

$$\left| \sum_{i=1}^q f(x_{\sigma_j(i)}) \right| \leq 8n \log m$$

for any $q = 1, 2, \dots, m$ and any $j = 1, 2, \dots, n$.

We construct our agreeable subset as follows. We include in our subset all items x_i such that $f(x_i) = 1$, as well as the $\lceil 4n \log m \rceil$ most preferred items of each agent that are not yet included. (If some agent has fewer than $\lceil 4n \log m \rceil$ items that are not yet included, we simply include all of that agent's items.) For $i = 1, 2, \dots, m$, let X_i be an indicator variable such that $X_i = 1$ if item x_i is included in the set and $X_i = -1$ if not. For any agent j and any $i = 1, 2, \dots, m$, we have

$$X_{\sigma_j(1)} + \dots + X_{\sigma_j(i)} \geq \min\{i, -8n \log m + 2 \cdot \lceil 4n \log m \rceil\} \geq 0.$$

By Proposition 1, this implies that the chosen set is necessarily agreeable to all agents. Moreover, our subset includes at most

$$\frac{m}{2} + (n+1) \cdot \lceil 4n \log m \rceil = \frac{m}{2} + O(\log m)$$

items, as desired. □

Next, we address the tightness of the bound in Theorem 4. Bohus's result is known to be asymptotically tight for constant n : Newman et al. [43] constructed, for every m that is a power of three, an example of three permutations whose discrepancy is $\Omega(\log m)$. While upper bounds on the discrepancy of permutations can be easily turned into upper bounds on the size of necessarily agreeable sets as seen above, lower bounds are somewhat more delicate. Nevertheless, Newman et al.'s examples satisfy stronger conditions than merely having a large discrepancy. One of these conditions, which we state in the following lemma, will be sufficient for proving a lower bound on the size of necessarily agreeable sets. The lemma is a restatement of Corollary 2 in the work of Newman et al. [43].

Lemma 3 ([43]). *Given any positive integer k , let $m = 3^k$ and $M = \{1, 2, \dots, m\}$. There exist three permutations $\sigma_1, \sigma_2, \sigma_3$ of M such that for any function $f : M \rightarrow \{-1, 1\}$, if $\Delta := \sum_{i \in M} f(i) \geq 1$, then there exist $1 \leq q \leq m$ and $1 \leq j \leq 3$ such that*

$$\sum_{i=1}^q f(\sigma_j(i)) \leq \frac{-k + 2\Delta - 2}{3}.$$

We now show that the bound in Theorem 4 is tight even when there are three agents. Recall that if there are two agents, it is possible to compute a subset of size $\lfloor \frac{m+2}{2} \rfloor$ that is necessarily agreeable to both agents (Theorem 2).

Theorem 5. *Suppose that $m = 3^k$ for some positive integer k . There exist preferences on single items of three agents such that every necessarily agreeable subset of items have size at least $m/2 + \Omega(\log m)$.*

Proof. Let $\sigma_1, \sigma_2, \sigma_3$ be the permutations of S from Lemma 3, where we use the set $S = \{x_1, x_2, \dots, x_m\}$ instead of the set $M = \{1, 2, \dots, m\}$. For each $j = 1, 2, 3$, let the preference on single items of agent j be $x_{\sigma_j(1)} \succeq_j^{\text{sing}} x_{\sigma_j(2)} \succeq_j^{\text{sing}} \dots \succeq_j^{\text{sing}} x_{\sigma_j(m)}$.

Consider any subset $T \subseteq S$ of size at most $m/2 + k/4$. We will show that T cannot be necessarily agreeable to all three agents, which immediately implies the theorem since $k = \log_3 m$. To see that this is the case, let $f_T : S \rightarrow \{-1, 1\}$ denote the indicator function of T , i.e., $f_T(x_i) = 1$ if $x_i \in T$ and $f_T(x_i) = -1$ if $x_i \notin T$. Since T is of size at most $m/2 + k/4$, we have $\Delta_T := \sum_{i=1}^m f_T(x_i) = |T| - |S \setminus T| \leq k/2$. If $\Delta_T < 0$, then T is trivially not necessarily agreeable, so we may assume that $\Delta_T \geq 0$. Since m is odd and Δ_T is an integer, we also have $\Delta_T \geq 1$. By Lemma 3, there exists $1 \leq q \leq m$ and $1 \leq j \leq 3$ such that

$$\sum_{i=1}^q f_T(\sigma_j(x_i)) \leq \frac{-k + 2\Delta_T - 2}{3} \leq \frac{-k + k - 2}{3} < 0.$$

By Proposition 1, T is not necessarily agreeable to agent j , as desired. \square

Theorems 4 and 5 show that the bound $m/2 + O(\log m)$ for the size of the smallest necessarily agreeable set is asymptotically tight. We next present a randomized algorithm

that, despite its simplicity, computes a necessarily agreeable subset of size $m/2 + O(\sqrt{m})$ in polynomial time. The algorithm works by first choosing whether to include each item independently with 50% probability, and then including the $O(\sqrt{m})$ most preferred items of each agent that were excluded in the first step.

For the analysis of the algorithm, we will require two probabilistic results. The first result, the Chernoff bound, is a fundamental result that gives an upper bound on the probability that a sum of independent random variables is far away from its expected value. For simplicity, we only state the bound for the case of symmetric ± 1 random variables, which suffices for our usage. This version can be found, for example, as Theorem A.1.1 in Alon and Spencer's book [4].

Lemma 4 (Chernoff bound). *Let X_1, X_2, \dots, X_r be i.i.d. random variables such that $\Pr[X_i = 1] = \Pr[X_i = -1] = 1/2$, and let $X := X_1 + \dots + X_r$. For any $a \geq 0$, we have*

$$\Pr[|X| \geq a] \leq \exp(-a^2/2r),$$

The second result, Lévy's inequality, yields a bound on the maximum of partial sums of independent random variables [35]. Once again, we state the inequality only for symmetric ± 1 random variables.

Lemma 5 (Lévy's inequality). *Let X_1, X_2, \dots, X_r be i.i.d. random variables such that $\Pr[X_i = 1] = \Pr[X_i = -1] = 1/2$, and let $Y_i := X_1 + \dots + X_i$ for $i = 1, 2, \dots, r$. For any real number x , we have*

$$\Pr\left[\max_{1 \leq i \leq r} |Y_i| \geq x\right] \leq 2 \Pr[|Y_r| \geq x].$$

We now present the algorithm and its analysis.

Theorem 6. *Assume that the number of agents is constant. Let $\epsilon \in (0, 1)$, and let $c > 0$ be a constant such that $e^{-c^2/2} \leq \epsilon/(2n)$. Consider the following randomized polynomial-time algorithm:*

1. *For each item, either include it in our set or not with probability 1/2, independently of the remaining items.*
2. *Include the $\lfloor c\sqrt{m} \rfloor$ most preferred items of each agent that were excluded in Step 1.*

With probability at least $1 - \epsilon$, the algorithm computes a subset of size $m/2 + O(\sqrt{m})$ that is necessarily agreeable to all agents.

Proof. Let X_1, X_2, \dots, X_m be independent random variables such that $X_i = +1$ if item x_i is included in our subset in the first step, and $X_i = -1$ if not. By the definition of the algorithm, each X_i is 1 or -1 with probability 1/2, independently of the other X_i 's.

For $j = 1, 2, \dots, m$, suppose that agent j ranks the single items as $x_{\sigma_j(1)} \succeq_j^{sing} x_{\sigma_j(2)} \succeq_j^{sing} \dots \succeq_j^{sing} x_{\sigma_j(m)}$. Let $Y_i^j := X_{\sigma_j(1)} + \dots + X_{\sigma_j(i)}$ for $i = 1, 2, \dots, n$. Applying Lemma 4 with $a = c\sqrt{m}$ yields

$$\Pr[|Y_m^j| \geq c\sqrt{m}] \leq e^{-c^2/2}.$$

Using Lemma 5 with the random variables $X_{\sigma_j(i)}$ for $i = 1, 2, \dots, n$, it follows that

$$\Pr \left[\max_{1 \leq i \leq m} |Y_i^j| \geq c\sqrt{m} \right] \leq 2e^{-c^2/2}.$$

Using the union bound over all agents j , we have

$$\Pr \left[\max_{1 \leq i \leq m} |Y_i^j| \geq c\sqrt{m} \text{ for some } j \in \{1, 2, \dots, m\} \right] \leq 2ne^{-c^2/2} \leq \epsilon.$$

Hence, with probability at least $1 - \epsilon$, $Y_i^j \in [-c\sqrt{m}, c\sqrt{m}]$ for all $i = 1, \dots, n$ and $j = 1, \dots, m$. Now, since in Step 2 we include the $\lfloor c\sqrt{m} \rfloor$ most preferred items of each agent that were excluded in Step 1, if we update the random variables to reflect these changes, we have $Y_i^j \geq 0$ for all i, j . By Proposition 1, the set that the algorithm returns is necessarily agreeable to all agents. Since we include at most $m/2 + c\sqrt{m}$ items in Step 1 and $c\sqrt{m}$ items for each agent in Step 2, the set contains at most

$$\frac{m}{2} + (n+1) \cdot c\sqrt{m} = \frac{m}{2} + O(\sqrt{m})$$

items, as desired. \square

4 Efficient Approximation

While our results in Sections 3 provide insights on small agreeable sets and how to compute them efficiently, an important issue is still left unaddressed by these results. In many instances, the minimum size of an agreeable set is much smaller than the worst-case bound over all instances with that number of agents and items. Indeed, an extreme example is when there is a single item that every agent likes more than all of the remaining items combined. In this case, it suffices to select that item alone. This results in a much smaller set than the worst-case bound, which is at least half of the items for any number of agents.

In this section, we investigate the problem of computing an agreeable subset of optimal size for any given instance, as opposed to one whose size matches the worst-case bound over all instances with the same number of agents and items. We show that finding an optimal agreeable set is computationally hard, and therefore focus on finding an approximate solution. We do so using two well-known models for representing preferences, namely the value oracle model and additive valuations. For each of these models, we present polynomial-time algorithms for computing an agreeable set of approximately optimal size. Moreover, we show that the approximation ratios obtained by our algorithms are asymptotically tight for both models.

4.1 General Preferences

We begin with a model in which agents can have arbitrary preferences on subsets of items. Recall that our results so far do not yield any guarantee on the approximation

ratio beyond the obvious $O(m)$ upper bound for arbitrary preferences over subsets of items. The goal of this section is to explore the approximation ratios that we can achieve in this general setting.

Before we move on to our results, let us be more precise about the model that we work with. First, we work with the agents' utility functions u_1, u_2, \dots, u_n instead of directly with the preferences themselves. Since the number of subsets of S is exponentially large, the utility functions take exponential space to write down. For this reason, it is undesirable to include them as part of the input. Instead, we work with the *value oracle model* [22], in which the algorithm can query the value of $u_i(T)$ for any subset $T \subseteq S$ and any $i = 1, 2, \dots, n$. We also note that we do not assume responsiveness of the agents' preferences in this section.

Our first result is a simple polynomial-time approximation algorithm with approximation ratio $O(m/\log m)$. Even though this approximation guarantee is only $\Omega(\log m)$ better than the obvious $O(m)$ bound, we will see later that this is already the best we can hope for in polynomial time.

Theorem 7. *There exists a polynomial-time $O(m/\log m)$ -approximation algorithm for computing a minimum size agreeable set in the value oracle model.*

Proof. We start by partitioning the set S of items into $\lceil \log m \rceil$ parts $S_1, \dots, S_{\lceil \log m \rceil}$, where each part is of size at most $\lceil m/\log m \rceil$. For each set $A \subseteq \{1, 2, \dots, \lceil \log m \rceil\}$, we check whether the set $\bigcup_{i \in A} S_i$ is agreeable or not by comparing each agent's value for the set to that for its complement. We then output the smallest agreeable set that we find. Since the number of possible sets A is linear in m , the running time of our algorithm is polynomial in m and n .

To prove the approximation guarantee of the algorithm, let S^* be a smallest agreeable set. Suppose that $|S^*| = k$. By monotonicity, the union of all sets S_i containing elements of S^* is also agreeable, and it is one of the sets that we check. Moreover, this union has size at most $k \cdot \lceil m/\log m \rceil$, implying that our algorithm indeed has approximation ratio $O(m/\log m)$. \square

Even though our algorithm is very simple, we show next that its approximation guarantee is in fact the best one can hope for, even when there is a single agent.

Theorem 8. *For every constant $c > 0$, there exists m_0 such that for every $m > m_0$, there is no (possibly randomized and adaptive) algorithm that makes at most $m^{c/8}$ queries to the value oracle and always outputs an agreeable set with expected size at most $m/(c \log m)$ times the optimum, even when there is only one agent.*

In other words, the above theorem implies that there is no polynomial time algorithm with approximation ratio $o(m/\log m)$. We note here that our lower bound is information-theoretic and is not based on any computational complexity assumptions. Moreover, it rules out any algorithm that makes a polynomial number of queries, not only those that run in polynomial time.

Proof of Theorem 8. Let $g : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ be a function such that

$$g(T) = \begin{cases} 1 & \text{if } |T| \geq m/2; \\ 0 & \text{otherwise.} \end{cases}$$

Moreover, for each subset $T^* \subseteq S$, let $f_{T^*} : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ denote the function

$$f_{T^*}(T) = \begin{cases} 1 & \text{if } |T| \geq m/2 \text{ or } T^* \subseteq T; \\ 0 & \text{otherwise.} \end{cases}$$

That is, f_{T^*} is g together with a planted solution T^* .

Consider any algorithm \mathcal{A} that makes at most $m^{c/8}$ queries. Assume for the moment that \mathcal{A} is deterministic. Let us examine a run of \mathcal{A} when the agent's utility function is g . Suppose that \mathcal{A} 's queries to g are on the sets $T_1, T_2, \dots, T_{\lfloor m^{c/8} \rfloor} \subseteq S$.

Let T^* be a subset of S of size $\lfloor c \log m/4 \rfloor$ chosen uniformly at random. Consider the queries that \mathcal{A} makes when the agent's utility function is f_{T^*} ; suppose that the queries made are on the sets $T'_1, T'_2, \dots, T'_{\lfloor m^{c/8} \rfloor} \subseteq S$. For each $j = 1, 2, \dots, \lfloor m^{c/8} \rfloor$, if $T_i = T'_i$ and $g(T_i) = f_{T^*}(T'_i)$ for all $i = 1, 2, \dots, j-1$, then \mathcal{A} goes through the same computation route for both g and f_{T^*} , and hence $T_j = T'_j$. Moreover, when both runs share the same computational route so far and $T_j = T'_j$, we can bound the probability that $g(T_j) \neq f_{T^*}(T'_j)$ as follows. First, if $|T_j| \geq m/2$, then $g(T_j)$ is always equal to $f_{T^*}(T'_j)$. Otherwise, we have

$$\Pr[g(T_j) \neq f_{T^*}(T'_j)] = \Pr[g(T_j) \neq f_{T^*}(T_j)] = \Pr[T^* \subseteq T_j].$$

If $|T_j| < |T^*|$, this probability is 0. Else, since T_j is independent of T^* , we can bound the probability as

$$\begin{aligned} \Pr[T^* \subseteq T_j] &= \frac{\binom{|T_j|}{\lfloor c \log m/4 \rfloor}}{\binom{m}{\lfloor c \log m/4 \rfloor}} \\ &= \left(\frac{|T_j|}{m} \right) \left(\frac{|T_j| - 1}{m - 1} \right) \cdots \left(\frac{|T_j| - \lfloor c \log m/4 \rfloor + 1}{m - \lfloor c \log m/4 \rfloor + 1} \right) \\ &\leq \left(\frac{|T_j|}{m} \right)^{\lfloor c \log m/4 \rfloor} \\ &\leq 2^{-\lfloor c \log m/4 \rfloor} \\ &\leq 2m^{-c/4}. \end{aligned}$$

By the union bound, the probability that the two sequences of queries are not identical is at most $(2m^{-c/4}) \cdot m^{c/8} = 2m^{-c/8}$, which is less than $1/2$ when m is sufficiently large. Furthermore, observe that when the two sequences are identical, \mathcal{A} must output an agreeable subset with respect to the utility function g ; any such set is of size at least $m/2$. Thus, the expected size of the output of \mathcal{A} when given the utility function f_{T^*} is more than $m/2 \cdot (1/2) = m/4$. However, the optimal agreeable set for f_{T^*} has size only

$\lfloor c \log m/4 \rfloor$. As a result, the expected size of the output of \mathcal{A} is more than $m/(c \log m)$ times the optimum, as desired.

Finally, note that if \mathcal{A} is randomized, we can use the above argument on each choice of randomness and average over all the choices, which gives a similar conclusion. \square

We remark that the same result holds even if we require the utility function of the agent to be subadditive or submodular.⁵ To obtain the proof for a subadditive utility function, for any $T \neq \emptyset$ such that $g(T) = 0$, we set instead $g(T) = 1/2$; we perform an analogous modification to f_{T^*} . Subadditivity holds for g since

$$g(A \cup B) \leq 1 = 1/2 + 1/2 \leq g(A) + g(B)$$

for any $A, B \neq \emptyset$, and similarly for f_{T^*} . The rest of the proof then proceeds as before.

On the other hand, more work is required to adapt the proof to submodular functions. In particular, we let $k = \lfloor c \log m/4 \rfloor$ and define g as follows:

$$g(T) = \begin{cases} 1 & \text{if } |T| \geq m/2, \\ 1 - \frac{1}{2^{|T|-k}(k+1)} & \text{if } k \leq |T| < m/2, \\ \frac{|T|}{k+1} & \text{otherwise.} \end{cases}$$

Likewise, for any T such that originally $f_{T^*}(T) = 0$, we modify the value of $f_{T^*}(T)$ to be the same as $g(T)$. One can check that g and f_{T^*} are submodular, and the proof again proceeds as before.

4.2 Additive Utilities

In this section, we assume that the agents' preferences are represented by additive utility functions. Each agent i has some nonnegative utility $u_i(x_j)$ for item x_j , and $u_i(T) = \sum_{x \in T} u_i(x)$ for any subset of items $T \subseteq S$.

Clearly, the problem of deciding whether there exists an agreeable set of a certain size is in NP. The following theorem shows that it is NP-complete, even when there are two agents. Note that if there is only one agent, it is easy to find an optimal agreeable set by repeatedly choosing an item that yields the highest utility to the agent among the remaining items until the set of chosen items is agreeable.

For the hardness proof, we will reduce from the following problem called BALANCED 2-PARTITION: Given a multiset A of non-negative integers, decide whether there exists a subset $B \subseteq A$ such that $|B| = |A \setminus B| = |A|/2$ and $\sum_{a \in B} a = \sum_{a \in A \setminus B} a = \sum_{a \in A} a/2$. Like the well-known 2-PARTITION problem where the cardinality constraint is not included, BALANCED 2-PARTITION is NP-hard, as shown in the following lemma.

Lemma 6. BALANCED 2-PARTITION is NP-hard.

⁵A function $f : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is said to be *submodular* if $f(A \cup \{x\}) - f(A) \geq f(B \cup \{x\}) - f(B)$ for any $A \subseteq B \subseteq \mathcal{S}$ and $x \in \mathcal{S} \setminus B$. Any submodular function is also subadditive.

Proof. We reduce from 2-PARTITION, a problem in which a multiset B of positive integers is given and the goal is to decide whether there exists a multiset $T \subseteq B$ such that $\sum_{b \in T} b = \sum_{b \in B \setminus T} b$. 2-PARTITION is known to be NP-complete (see, e.g., [24]).

Given a 2-PARTITION instance, we create a BALANCED 2-PARTITION instance as follows. Let A be the multiset containing all elements of B and $|B|$ additional zeros. Clearly, the reduction runs in polynomial time. We show that B is a YES instance of 2-PARTITION if and only if A is a YES instance of BALANCED 2-PARTITION.

(YES Case) Suppose that B is a YES instance of 2-PARTITION, i.e., there exists $T \subseteq B$ such that $\sum_{b \in T} b = \sum_{b \in B \setminus T} b$. Let $S \subseteq A$ be the multiset containing all elements of T and $|B| - |T|$ additional zeros. Clearly, $|S| = |B| = |A|/2$ and $\sum_{a \in S} a = \sum_{b \in T} b = \sum_{b \in B} b/2 = \sum_{a \in A} a/2$, meaning that A is a YES instance of BALANCED 2-PARTITION as desired.

(NO Case) We prove the contrapositive; suppose that A is a YES instance of BALANCED 2-PARTITION. This means that there exists $S \subseteq A$ of size $|A|/2 = |B|$ such that $\sum_{a \in S} a = \sum_{a \in A \setminus S} a$. Let T be the subset of B containing all elements of B whose corresponding elements are included in S . Clearly, we have $\sum_{b \in T} b = \sum_{a \in S} a = \sum_{a \in A \setminus S} a = \sum_{b \in B \setminus T} b$. Hence B is a YES instance of 2-PARTITION. \square

Theorem 9. *For two agents with additive utility functions, it is NP-hard to decide whether there is an agreeable set of size exactly $m/2$.*

Proof. The reduction from BALANCED 2-PARTITION proceeds as follows. Let $a_1, \dots, a_{|A|}$ be the elements of A . The set S contains $|A|$ items $x_1, \dots, x_{|A|}$, each associated with an element of A . The utility functions are then defined by $u_1(x_i) = a_i$ and $u_2(x_i) = M - a_i$, where $M = \sum_{a \in A} a$. We show next that this reduction is indeed a valid reduction.

(YES Case) Suppose that there exists $B \subseteq A$ such that $|B| = |A|/2$ and $\sum_{a \in B} a = \sum_{a \in A} a/2$. Let T be the set of all items corresponding to the elements of B . It is obvious that T has size $|A|/2 = m/2$ and that T is agreeable.

(NO Case) We prove the contrapositive; suppose that there is an agreeable subset $T \subseteq S$ of size $m/2$. Let $B \subseteq A$ be the set of elements corresponding to the items in T . Since T is agreeable, $\sum_{x \in T} u_i(x) \geq \sum_{x \in S \setminus T} u_i(x)$ for $i = 1, 2$. When $i = 1$, this implies that $\sum_{a \in B} a \geq \sum_{a \in A} a/2$. When $i = 2$, using the fact that $|T| = m/2$, we have $\sum_{a \in B} a \leq \sum_{a \in A} a/2$. It follows that $\sum_{a \in B} a = \sum_{a \in A} a/2$. Since $|B| = m/2 = |A|/2$, this concludes the proof. \square

Theorem 9 shows that the problem is weakly NP-hard even when there are two agents. Nevertheless, when the number of agents is constant, the following theorem shows that there exists a pseudo-polynomial time dynamic programming algorithm for computing an optimal agreeable set. In particular, the problem is not strongly NP-hard for any constant number of agents.

Theorem 10. *For any constant number of agents with additive utility functions, there exists a pseudo-polynomial time algorithm that computes an agreeable set of minimum size.*

Proof. The algorithm uses dynamic programming. Assume that the utilities of agent i for the items are nonnegative integers with sum σ_i . We construct a table Σ of size $(m+1)(\sigma_1+1)\dots(\sigma_n+1)$, where for each $0 \leq m' \leq m$ and each tuple (y_1, \dots, y_n) with $0 \leq y_i \leq \sigma_i$, the entry $\Sigma(m', y_1, \dots, y_n)$ of the table corresponds to the minimum number of items from among the items $x_1, x_2, \dots, x_{m'}$ that we need to include so that agent i has utility exactly y_i for all i (if this is achievable). Initially we have $\Sigma(0, 0, \dots, 0) = 0$ and $\Sigma(m', y_1, \dots, y_n) = \infty$ otherwise. We then iterate through the values of m' in increasing order. For each $m' \geq 1$, we update the entries of the table as follows:

- If $u_i(x_{m'}) \leq y_i$ for all i and

$$1 + \Sigma(m' - 1, y_1 - u_1(x_{m'}), \dots, y_n - u_n(x_{m'})) < \Sigma(m' - 1, y_1, \dots, y_n),$$

$$\text{let } \Sigma(m', y_1, \dots, y_n) = 1 + \Sigma(m' - 1, y_1 - u_1(x_{m'}), \dots, y_n - u_n(x_{m'})).$$

- Else, let $\Sigma(m', y_1, \dots, y_n) = \Sigma(m' - 1, y_1, \dots, y_n)$.

Finally, we look up the entries $\Sigma(m, y_1, \dots, y_n)$ such that $y_i \geq \sigma_i/2$ for all i and return the minimum value over all such entries. The algorithm runs in time $O(m\sigma_1 \dots \sigma_n)$. Note that if we also want to return an agreeable set (rather than just the size), we can also keep track of the sets of items along with the values in our table. \square

While there is a pseudo-polynomial time algorithm for the problem when the number of agents is constant, we show next that if the number of agents is not constant, the problem becomes strongly NP-hard. In other words, there is no pseudo-polynomial time algorithm for this variant unless $P=NP$.

Theorem 11. *When the number of agents is not constant, it is strongly NP-hard to decide whether there is an agreeable set of size exactly $(m+1)/2$.*

Proof. We reduce from 3SAT. Given a 3SAT formula ϕ with m' clauses $C_1, C_2, \dots, C_{m'}$ on n' variables $y_1, y_2, \dots, y_{n'}$, let there be $n = m' + n'$ agents, where we abuse notation and call the agents $C_1, C_2, \dots, C_{m'}, y_1, y_2, \dots, y_{n'}$, and $m = 2n' + 1$ items, where $2n'$ items correspond to all the literals $y_1, \neg y_1, y_2, \neg y_2, \dots, y_{n'}, \neg y_{n'}$ and the remaining item is called a . We assume without loss of generality that each clause of ϕ has at least two variables—it is obvious that every 3SAT formula can be transformed into this form in polynomial time. The utility functions of the agents are defined by

$$u_{C_i}(b) = \begin{cases} 1 & \text{if } b = a \text{ or the literal } b \text{ is present in } C_i; \\ 0 & \text{otherwise;} \end{cases}$$

and

$$u_{y_i}(b) = \begin{cases} 1 & \text{if } b = a, b = y_i, \text{ or } b = \neg y_i; \\ 0 & \text{otherwise.} \end{cases}$$

We show next that this is a valid reduction. First, note that all of the integer parameters are polynomial in the size of the input. Hence, we are left to show that YES and NO instances of 3SAT map to YES and NO instances of our problem respectively.

(YES Case) Suppose that there exists an assignment that satisfies ϕ . For each y_i , let b_i be the literal of y_i that is true according to this assignment. Let $T = \{a, b_1, b_2, \dots, b_{n'}\}$. Since each clause C_j is satisfied by the assignment, we have $\sum_{i=1}^{n'} u_{C_j}(b_i) \geq 1$. It follows that $\sum_{x \in T} u_{C_j}(x) \geq 2$, and therefore $T \succeq_{C_j} S \setminus T$. Moreover, for each variable y_i , we have $\sum_{x \in T} u_{y_i}(x) = 2$, which also implies that $T \succeq_{y_i} S \setminus T$. As a result, T is an agreeable set of size $n' + 1 = (m + 1)/2$ as desired.

(NO Case) We prove the contrapositive; suppose that there exists an agreeable set $T \subseteq S$ of size $(m + 1)/2 = n' + 1$. We assume without loss of generality that $a \in T$; indeed, since the utility of any agent for a is at least as much as the utility of the agent for any other item, if $a \notin T$ we can replace an arbitrary item in T by a and maintain the agreeability of T .

Since $T \succeq_{y_i} S \setminus T$, at least one literal corresponding to y_i is included in T . Moreover, since the size of T is $n' + 1$ and $a \in T$, exactly one literal of each y_i is in T ; let b_i be this literal. Consider the assignment to the variables such that all the b_i 's are set to true. Since $T \succeq_{C_j} S \setminus T$ for every C_j and C_j contains at least two literals, at least one literal in C_j is set to true by this assignment. Hence the assignment satisfies the formula ϕ . \square

Given that computing an agreeable set of minimum size is NP-hard, it is natural to attempt to find an approximation algorithm for the problem. When the utilities are additive, this turns out to be closely related to approximating the classical problem SET COVER. In SET COVER, we are given a ground set U and a collection \mathcal{C} of subsets of U . The goal is to select a minimum number of subsets whose union is the entire set U .

SET COVER was one of the first problems shown to be NP-hard in Karp's seminal paper [28]. Since then, its approximability has been intensively studied and is now well understood. A simple greedy algorithm yields a $(\ln |U| + 1)$ -approximation for the problem [27, 36]. On the other hand, a long line of work in hardness of approximation [3, 20, 38, 42, 45] culminates in Dinur and Steurer's work [19], in which the NP-hardness of approximating SET COVER within a factor of $(1 - \varepsilon) \ln |U|$ was proved for every constant $\varepsilon > 0$.

The first connection we will make between SET COVER and approximating minimum size agreeable set is on the negative side—we will show that any inapproximability result for SET COVER can be translated to that for approximating minimum size agreeable set as well. To do so, we will first state Dinur and Steurer's result more precisely.

Lemma 7 ([19]). *For every constant $\varepsilon > 0$, there is a polynomial time reduction from any 3SAT formula ϕ to a SET COVER instance (U, \mathcal{C}) and a function $f(U)$ which is polynomial in $|U|$ such that*

- (Completeness) *if ϕ is satisfiable, the optimum of (U, \mathcal{C}) is at most $f(U)$;*
- (Soundness) *if ϕ is unsatisfiable, the optimum of (U, \mathcal{C}) is at least $((1 - \varepsilon) \ln |U|) f(U)$.*

We are now ready to prove the hardness of approximating minimum size agreeable set.

Theorem 12. *For any constant $\delta > 0$, it is NP-hard to approximate minimum size agreeable set to within a factor $(1 - \delta) \ln n$ of the optimum.*

Proof. Let $\varepsilon = \delta/2$. Given a 3SAT formula ϕ , we first use Dinur and Steurer's reduction to produce a SET COVER instance (U, \mathcal{C}) . Let there be $|U|$ agents, each of whom is associated with a distinct element of U ; it is convenient to think of the set of agents as simply U . As for the items, let there be one item for each subset $C \in \mathcal{C}$ and additionally let there be one special item called t . In other words, $S = \mathcal{C} \cup \{t\}$.

The utility function of each agent $a \in U$ is then defined by

$$u_a(s) = \begin{cases} |\{C \in \mathcal{C} \mid a \in C\}| - 1 & \text{if } s = t; \\ 1 & \text{if } s \in \mathcal{C} \text{ and } a \in s; \\ 0 & \text{otherwise.} \end{cases}$$

We show next that this reduction indeed gives the desired inapproximability result.

(Completeness) If ϕ is satisfiable, then there are $f(U)$ subsets from \mathcal{C} that together cover U . We can take T to contain all of these subsets and the special item t . Clearly, T has size $f(U) + 1$ and is agreeable.

(Soundness) If ϕ is unsatisfiable, then any set cover of (S, \mathcal{C}) contains at least $((1 - \varepsilon) \ln |U|)f(U)$ subsets. Consider any agreeable set T . For each $a \in U$, from our definition of $u_a(t)$, the set T must include at least one subset that contains a . In other words, $T \setminus \{t\}$ is a set cover of (S, \mathcal{C}) . Hence, $|T| \geq ((1 - \varepsilon) \ln |U|)f(U)$.

The two parts together imply that it is NP-hard to approximate minimum size agreeable set to within a factor $\frac{((1 - \varepsilon) \ln |U|)f(U)}{f(U) + 1}$ of the optimum. This ratio is at least $(1 - \delta) \ln n$ when $f(U) \geq 2/\delta$, which can be assumed without loss of generality (since otherwise we can solve the SET COVER instance in time $|U|^{O(f(U))} = |U|^{O(1)}$, implying that P = NP). \square

Unlike the above inapproximability result, it is unclear how algorithms for SET COVER can be used to approximate minimum size agreeable set. Fortunately, our problem is in fact a special case of a generalization of SET COVER called COVERING INTEGER PROGRAM (CIP), which can be written as follows:

$$\begin{aligned} & \text{minimize } c^T x \\ & \text{subject to } Ax \geq 1, \\ & \quad 0 \leq x \leq u, \\ & \quad x \in \mathbb{Z}^m, \end{aligned}$$

where $c, u \in \mathbb{R}^m$ and $A \in \mathbb{R}^{n \times m}$ are given as input.

The problem of finding a minimum size agreeable set can be formulated in this form by setting c, u and A as follows:

$$c_s = 1 \qquad \forall s \in S$$

$$\begin{aligned}
u_s &= 1 & \forall s \in S \\
A_{i,s} &= \frac{2u_i(s)}{\sum_{s' \in S} u_i(s')} & \forall i = 1, 2, \dots, n, \forall s \in S
\end{aligned}$$

Similarly to SET COVER, the approximability of CIP is well studied. In particular, the problem is known to be approximable to within a factor $O(\log n)$ of the optimum in polynomial time [32]. This immediately implies an $O(\log n)$ -approximation algorithm for finding a minimum size agreeable set as well.

Theorem 13. *For agents with additive utility functions, there exists a polynomial-time $O(\log n)$ -approximation algorithm for computing a minimum size agreeable set.*

5 Discussion

In this paper, we introduce the notion of agreeability, which captures an agent’s acceptance of the set of items allocated to her group, and present a number of fundamental results on the notion. For any number of agents and items, we derive a tight upper bound on the number of items that may need to be included in an agreeable subset. We also present polynomial-time algorithms for computing an agreeable set whose size matches the upper bound or approximates the optimal size for a given instance using well-known models for representing preferences. To the best of our knowledge, this is the first work that studies a model in which a single group of agents collectively receive a subset of items.

Our work suggests a number of possible future directions. With polynomial-time algorithms for computing an agreeable set whose size matches the upper bound for two and three agents in hand, a natural question is whether we can similarly obtain efficient algorithms when there are more agents. The algorithm for three agents is already quite involved, so one might suspect that the problem is intractable for larger numbers of agents. If that were to be the case, it would be useful to have a confirmation by means of a hardness result, even for some fixed large number of agents. Since the problem is a search problem for which we know that a solution always exists, it cannot be NP-hard, but could potentially be hard with respect to a subclass of TFNP such as PPAD or PLS. One could also investigate the complexity of deciding the existence of agreeable subsets of certain sizes for which there is no guarantee of existence, as we do in Theorems 9 and 11.

Another avenue for future work is to extend the notion of agreeability to more general settings. For instance, a motivating example that we give is that the group of agents receive some items as prizes from a team competition that they won against another group. One could consider a generalization where there are more than two competing groups. However, in this case there are several reasonable ways of defining agreeability, since we do not know how the remaining items are distributed among the remaining groups. One possibility is to require that each agent in the group finds the set of items to be worth at least $1/k$ of the whole set, where k is the number of groups. An alternative definition is to impose the condition that for each agent in the group, there is a way to

partition the remaining items among the other $k - 1$ groups so that the agent does not envy any of the other groups. While both definitions reduce to our notion of agreeability in the case of two groups and additive utilities, the equivalence ceases to hold when there are at least three groups or if utilities are not additive. As such, the results that we can obtain are likely to depend on the definition that we use.

An interesting related question that goes beyond our setting is the following: When is it possible to obtain a fair allocation between two or more groups of agents, where, like in our work, agents in the same group share the same set of items? As mentioned in the introduction, this question has been a subject of recent attention, but much work remains to be done. For example, in the traditional fair division setting with one agent per group, the undercut procedure [16] is an elegant method for computing an envy-free allocation whenever one exists. Whether we can come up with a similar procedure for the group setting is an appealing direction that we leave for future research.

Acknowledgments

The authors would like to thank Nikhil Bansal and Aleksandar Nikolov for helpful discussions, and the reviewers of the 25th International Joint Conference on Artificial Intelligence (New York, USA, July 2016), the 26th International Joint Conference on Artificial Intelligence (Melbourne, Australia, August 2017), and Artificial Intelligence Journal for many insightful comments. Warut Suksompong acknowledges support from a Stanford Graduate Fellowship.

References

- [1] David S. Ahn and Santiago Oliveros. Combinatorial voting. *Econometrica*, 80(1):89–141, 2012.
- [2] Noga Alon. Splitting necklaces. *Advances in Mathematics*, 63(3):247–253, 1987.
- [3] Noga Alon, Dana Moshkovitz, and Shmuel Safra. Algorithmic construction of sets for k -restrictions. *ACM Transactions on Algorithms*, 2(2):153–177, 2006.
- [4] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. Wiley Publishing, 3rd edition, 2008.
- [5] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms*, 13(4):52, 2017.
- [6] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.

- [7] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.
- [8] Imre Bárány. A short proof of Kneser’s conjecture. *Journal of Combinatorial Theory, Series A*, 25(3):325–326, 1978.
- [9] Salvador Barberà, Walter Bossert, and Prasanta K. Pattanaik. Ranking sets of objects. In Salvador Barbera, Peter Hammond, and Christian Seidl, editors, *Handbook of Utility Theory, Volume 2*, pages 893–977. Kluwer Academic Publishers, 2004.
- [10] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–709, 2011.
- [11] Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 91–97, 2018.
- [12] Géza Bohus. On the discrepancy of 3 permutations. *Random Structures and Algorithms*, 1(2):215–220, 1990.
- [13] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Fair division under ordinal preferences: computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence*, pages 387–392, 2010.
- [14] Sylvain Bouveret and Jérôme Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research*, 32:525–564, 2008.
- [15] Sylvain Bouveret and Michel Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, 2016.
- [16] Steven J. Brams, D. Marc Kilgour, and Christian Klamler. The undercut procedure: an algorithm for the envy-free division of indivisible items. *Social Choice and Welfare*, 39(2–3):615–631, 2012.
- [17] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 305–322, 2016.
- [18] John P. Dickerson, Jonathan Goldman, Jeremy Karp, Ariel D. Procaccia, and Thomas Sandholm. The computational rise and fall of fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 1405–1411, 2014.
- [19] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Proceedings of the 46th Symposium on Theory of Computing*, pages 624–633, 2014.

- [20] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [21] Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009.
- [22] Uriel Feige, Vahab S. Mirrokni, and Jan Vondrák. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- [23] Duncan K. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7(1):45–98, 1967.
- [24] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [25] Laurent Gourvès. Agreeable sets with matroidal constraints. *Journal of Combinatorial Optimization*, 2018. Forthcoming.
- [26] Joshua E. Greene. A new short proof of Kneser’s conjecture. *American Mathematical Monthly*, 109(10):918–920, 2002.
- [27] David S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9(3):256–278, 1974.
- [28] Richard M. Karp. Reducibility among combinatorial problems. In *Proceedings of Complexity of Computer Computations*, pages 85–103, 1972.
- [29] Ian Kash, Ariel D. Procaccia, and Nisarg Shah. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, (51):579–603, 2014.
- [30] Martin Kneser. Aufgabe 360. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 58(2):27, 1955.
- [31] David A. Kohler and R. Chandrasekaran. A class of sequential games. *Operations Research*, 19(2):270–277, 1971.
- [32] Stavros G. Kolliopoulos and Neal E. Young. Approximation algorithms for covering/packing integer programs. *Journal of Computer and System Sciences*, 71(4):495–505, 2005.
- [33] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 523–529, 2016.
- [34] Jérôme Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1):37–71, 2004.
- [35] Paul Lévy. *Théorie de l’addition des variables aléatoires*. Gauthier-Villiers, 1937.

- [36] László Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383–390, 1975.
- [37] László Lovász. Kneser’s conjecture, chromatic number and homotopy. *Journal of Combinatorial Theory, Series A*, 25(3):319–324, 1978.
- [38] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.
- [39] Pasin Manurangsi and Warut Suksompong. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences*, 89:100–108, 2017.
- [40] Pasin Manurangsi and Warut Suksompong. Computing an approximately optimal agreeable set of items. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 338–344, 2017.
- [41] Jiří Matoušek. A combinatorial proof of Kneser’s conjecture. *Combinatorica*, 24(1):163–170, 2004.
- [42] Dana Moshkovitz. The projection games conjecture and the NP-hardness of $\ln n$ -approximating set-cover. *Theory of Computing*, 11:221–235, 2015.
- [43] Alantha Newman, Ofer Neiman, and Aleksandar Nikolov. Beck’s three permutations conjecture: a counterexample and some consequences. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 253–262, 2012.
- [44] Kirk Pruhs and Gerhard J. Woeginger. Divorcing made easy. In *Proceedings of the 6th International Conference on Fun with Algorithms*, pages 305–314, 2012.
- [45] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 475–484, 1997.
- [46] Erel Segal-Halevi and Shmuel Nitzan. Proportional cake-cutting among families. *CoRR*, abs/1510.03903, 2015.
- [47] Erel Segal-Halevi and Shmuel Nitzan. Envy-free cake-cutting among families. *CoRR*, abs/1607.01517, 2016.
- [48] Erel Segal-Halevi and Warut Suksompong. Democratic fair allocation of indivisible goods. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 482–488, 2018.
- [49] Forest W. Simmons and Francis Edward Su. Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Mathematical Social Sciences*, 45(1):15–25, 2003.
- [50] Piotr Skowron, Piotr Faliszewski, and Jérôme Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216, 2016.

- [51] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.
- [52] Warut Suksompong. Assigning a small agreeable set of indivisible items to multiple players. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 489–495, 2016.
- [53] Warut Suksompong. Approximate maximin shares for groups of agents. *Mathematical Social Sciences*, 92:40–47, 2018.
- [54] Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.
- [55] Lirong Xia and Jérôme Lang. A dichotomy theorem on the existence of efficient or neutral sequential voting correspondences. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 342–347, 2009.