

# Group Fairness in Committee Selection

YU CHENG, Computer Science, Duke University

ZHIHAO JIANG, IIS, Tsinghua University

KAMESH MUNAGALA, Computer Science, Duke University

KANGNING WANG, Computer Science, Duke University

In this paper, we study *fairness* in committee selection problems. We consider a general notion of fairness via stability: A committee is *stable* if no coalition of voters can deviate and choose a committee of proportional size, so that all these voters strictly prefer the new committee to the existing one. Our main contribution is to extend this definition to stability of a distribution (or lottery) over committees. We consider two canonical voter preference models: the APPROVAL SET setting where each voter approves a set of candidates and prefers committees with larger intersection with this set; and the RANKING setting where each voter ranks committees based on how much she likes her favorite candidate in a committee. Our main result is to show that stable lotteries always exist for these canonical preference models. Interestingly, given preferences of voters over committees, the procedure for computing an approximately stable lottery is the *same* for both models and therefore extends to the setting where some voters have the former preference structure and others have the latter. Our existence proof uses the probabilistic method and a new large deviation inequality that may be of independent interest.

CCS Concepts: • **Theory of computation** → *Algorithmic game theory; Solution concepts in game theory;* • **Mathematics of computing** → *Combinatoric problems; Probabilistic algorithms.*

Additional Key Words and Phrases: fairness, committee selection, social choice

## ACM Reference Format:

Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2019. Group Fairness in Committee Selection. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19)*, June 24–28, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3328526.3329577>

## 1 INTRODUCTION

One of the central problems in social choice theory is *committee selection*, or *multi-winner elections*. In this problem, there is a set of voters (or agents)  $\mathcal{N} = [n] = \{1, \dots, n\}$ , and a set of candidates  $C = [m]$ . A committee is a subset of candidates, and the goal is to find a committee of given size  $K$ . Committee selection arises in choosing a parliament, in making group hiring decisions, and in computer systems [13, 16]. A long line of recent literature [2, 3, 17, 18, 22, 23] has studied the complexity and axiomatization of voting rules in this setting.

One classic objective in committee selection is *fairness* or *proportionality*: Every demographic of voters should feel that they have been fairly represented. They should not have the incentive to deviate and choose their own committee of proportionally smaller size which gives all of them higher utility. In the typical setting where these demographic slices are not known upfront, the notion of proportionality attempts to be fair to all subsets of voters. This general idea dates back more than a century [7], and has recently received significant attention [2, 3, 5, 6, 19, 23]. In fact,

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329577>

there are several elections, both at a group level and national level, that attempt to find committees (or parliaments) that provide approximately proportional representation. For instance, the popular Single Transferable Vote (STV) rule is used in parliamentary elections in Ireland and Australia, and in several municipal elections in the USA. This rule attempts to find a proportional solution.

Fairness in committee selection arises in many other applications outside of social choice as well. For example, consider a shared cache for data items in a multi-tenant cloud system, where each data item is used by several long-running applications [13, 16]. Each data item can be treated as a candidate, and each application as a voter whose utility for an item corresponds to the speedup obtained by caching that item. In this context, we need a *fair* caching policy that provides proportional speedup to all applications.

In this paper, we propose a new notion of proportionality in committee selection that generalizes several previously considered notions. Our main contribution is to show that stable solutions always exist for the stability notion that we propose, and such solutions can be computed efficiently. In contrast, stable solutions may not exist for some of the previously studied notions, and for the notions where stable solutions do exist, we do not know how to compute them efficiently (see Section 1.5 for a detailed discussion).

## 1.1 Preference Models

Before proceeding further, we define the preference model of the voters for committees. We consider two canonical ordinal preference models over committees, both of which have been extensively studied in social choice literature.

**APPROVAL SET.** In this model [2, 3, 5, 23], each voter  $v$  specifies an *approval set*  $A_v \subseteq C$  of candidates. Given two committees  $S_1$  and  $S_2$ ,  $S_1 \succ_v S_2$  iff  $|S_1 \cap A_v| > |S_2 \cap A_v|$ , i.e., the voter strictly prefers committees in which she has more approved candidates.

**RANKING.** In this model [8], each voter  $v$  has a preference order over candidates in  $C$ . In this case,  $S_1 \succ_v S_2$  iff  $v$ 's favorite candidate in  $S_1$  is ranked higher (in her preference ordering) than her favorite candidate in  $S_2$ . The Chamberlin-Courant voting rule Chamberlin and Courant [6] for committee selection finds the social optimum assuming a cardinal preference function of this form.

These models have been extensively studied because it is relatively simple to elicit an approval vote or a ranking over candidates. Viewed in terms of underlying cardinal utility functions<sup>1</sup>, both these models are special cases of *submodular utilities*: Voter  $v$  has utility  $u_v(S)$  for committee  $S \subseteq C$ , where  $u_v$  is a submodular set function. The voter prefers committees that give her larger utility. The APPROVAL SET case corresponds to  $u_v(S) = f_v(|S \cap A_v|)$  for some increasing, concave function  $f_v$ . The RANKING case can be modeled as follows: Voter  $v$  has utility  $u_{vk}$  for candidate  $k \in C$ , and we set  $u_v(S) = \max_{k \in S} u_{vk}$ . The Chamberlin-Courant rule sets  $u_{vk} = m - \ell$  (i.e., the Borda score) if candidate  $k$  is ranked at position  $\ell$  by voter  $v$ .

## 1.2 Group Fairness via Stability

The notion of fairness we study is defined for arbitrary ordinal preferences over committees. We study fairness via the notion of stability, which has been extensively studied (in similar or more restricted forms) in economics [12, 20, 24] and computer science [2, 3, 9, 10]. We first define the notion of *capture count*.

<sup>1</sup>Note that we do not elicit these cardinal utilities.

DEFINITION 1 (CAPTURE COUNT). Given two committees  $S_1, S_2 \subseteq C$ , the *capture count* of  $S_2$  over  $S_1$  is the number of voters who strictly prefer  $S_2$  to  $S_1$ :

$$V(S_1, S_2) = |\{v \in \mathcal{N} \mid S_2 \succ_v S_1\}|.$$

We are now ready to define the notion of stability:

DEFINITION 2 (STABLE COMMITTEES). Given a committee  $S \subseteq C$  of size  $K$ , we say that a committee  $S' \subseteq C$  blocks  $S$  iff

$$V(S, S') \geq \frac{|S'|}{K} \cdot n.$$

For any  $1 \leq L \leq K$ , a committee  $S$  is *L-stable* if there are no committees  $S'$  of size at most  $L$  that block it. We term  $K$ -stable as just “stable”.

In other words, for any sub-group of voters of size  $\beta n$  for some  $\beta \in (0, 1]$ , the committee  $S$  is fair in the sense that this sub-group cannot form another committee of size  $\lfloor \beta K \rfloor$  so that all members of this sub-group are strictly better off. The notion of  $L$ -stability restricts this definition to “small” subgroups with  $\beta \leq \frac{L}{K}$ .

The classical interpretation of this concept in Economics [12, 20, 24] is in terms of *taxation*: Each voter has a dollar of money. If the goal is to find a committee of size  $K$ , we assume each committee member has a cost of  $\frac{n}{K}$ . A committee  $S$  of size  $K$  is stable if there is no subset of voters who can deviate and “buy” another committee  $S'$  using their share of money, so that all voters in this deviating coalition are strictly better off with  $S'$  than with  $S$ . Note that the amount of money required to buy  $S'$  is precisely  $\frac{|S'|}{K} \cdot n$ , and the amount of money available with the deviating coalition of voters who strictly prefer  $S'$  to  $S$  is precisely the capture count  $V(S, S')$ . Therefore  $S'$  is blocking if the deviating coalition of voters have sufficient funds to pay for  $S'$ .

Note that a stable outcome  $S$  is (weakly) Pareto-optimal among committees of size at most  $K$ ; if it were not, consider a Pareto-dominating committee  $S'$ . Since  $V(S, S') = n$ , the committee  $S'$  would be blocking, which contradicts the stability of  $S$ . In fact, it says something stronger: For every coalition of voters, a stable committee is also Pareto-optimal relative to committees whose size is suitably scaled down. Stability also logically implies other notions of stability considered in multi-winner election literature. For instance, the special case of 1-stability has been extensively studied as *justified representation*, and we discuss these connections further in Section 1.5.

In the RANKING setting, we can make the following simple observation:  $V(S, S') \leq \sum_{j \in S'} V(S, j)$ , since each deviating voter  $v$  prefers  $S'$  because of her favorite candidate in it. Therefore, by an averaging argument, if  $S'$  blocks  $S$ , then so does at least one  $j \in S'$ . This immediately gives the following observation:

OBSERVATION 1. *In the RANKING setting, a committee is stable iff it is 1-stable.*

### 1.3 Stable Lotteries and Approximate Stability

Given the strength of the definition, it is no surprise that stable committees may not exist for the RANKING setting. Consider the following example with cyclic preferences:

*Example 1.1.* There are  $n = 6$  voters  $\{1, 2, 3, 4, 5, 6\}$  and  $m = 6$  candidates  $\{a, b, c, d, e, f\}$ . We need to choose a committee of size  $K = 3$ . The rankings of the voters are as follows:

Voter	Preferences
1	$a > b > c > d \geq e \geq f$
2	$b > c > a > d \geq e \geq f$
3	$c > a > b > d \geq e \geq f$
4	$d > e > f > a \geq b \geq c$
5	$e > f > d > a \geq b \geq c$
6	$f > d > e > a \geq b \geq c$

Note that any committee chooses either at most one candidate from  $\{a, b, c\}$  or at most one candidate from  $\{d, e, f\}$ . Assume w.l.o.g. that the former happens, and candidate  $a$  is chosen (or no one in  $\{a, b, c\}$  is chosen). Then voters  $\{2, 3\}$  can deviate and choose candidate  $c$ . Therefore, there is no stable committee on this example.

In light of this impossibility result, we extend Definition 2 to allow randomization in choosing the committee. Given a committee size  $K$ , we let  $\Delta$  denote a distribution (or lottery) over committees of size  $K$ . Our first contribution is the following definition of stable lotteries.

**DEFINITION 3 (STABLE LOTTERIES).** A distribution (or lottery)  $\Delta$  over committees of size  $K$  is said to be  $L$ -stable ( $1 \leq L \leq K$ ) iff for all committees  $S' \subseteq C$  with  $|S'| \leq L$ , we have:

$$\mathbf{E}_{S \sim \Delta} [V(S, S')] < \frac{|S'|}{K} \cdot n .$$

We term  $K$ -stable lotteries as just “stable”.

In the taxation interpretation, the above definition says the following: For any committee  $S'$ , if in every realization  $S$  of  $\Delta$ , the voters who strictly prefer  $S'$  over  $S$  pay for  $S'$ , then in expectation there is not enough money to pay for  $S'$ . This justifies why the voters do not deviate to  $S'$  given  $\Delta$ . Further, we note that implementing a lottery is feasible in several computer systems applications. For instance, consider for example caching data items that are shared by several applications discussed above. In this context, a lottery over possible cache allocations can be implemented by splitting time into chunks, and caching the allocations in the lottery in different chunks of time.

We note for the RANKING setting, Observation 1 extends to lotteries because

$$\mathbf{E}_{S \sim \Delta} [V(S, S')] \leq \sum_{j \in S'} \mathbf{E}_{S \sim \Delta} [V(S, j)] .$$

**OBSERVATION 2.** *In the RANKING setting, a lottery is stable iff it is 1-stable.*

*Approximate Stability.* We next define approximately stable lotteries; this notion will play an important role in our computational results.

**DEFINITION 4 ( $\epsilon$ -APPROXIMATE STABILITY).** For any  $\epsilon > 0$ , a distribution  $\Delta$  over committees of size  $K$  is said to be  $\epsilon$ -approximately  $L$ -stable if for all committees  $S' \subseteq C$  with  $|S'| \leq L$ , we have:

$$\mathbf{E}_{S \sim \Delta} [V(S, S')] \leq (1 + \epsilon) \frac{|S'|}{K} \cdot n .$$

We term  $\epsilon$ -approximately  $K$ -stable lotteries as just  $\epsilon$ -approximately stable.

## 1.4 Our Results

Our main result in this paper is the following theorem, which states that a stable lottery always exists and it is reasonably tractable. Recall that (as in Definition 3) a lottery is  $L$ -stable if there are no committees of size *at most*  $L$  that block it, and stable lotteries refer to  $K$ -stable lotteries.

**THEOREM 1.2 (MAIN).** *In both the APPROVAL SET and RANKING settings with  $m$  candidates, for any committee size  $K$ , a stable lottery over committees of size  $K$  always exists. Furthermore, for any  $\epsilon > 0$ , an  $\epsilon$ -approximately  $L$ -stable lottery can be computed in  $\text{poly}(m^L, \frac{1}{\epsilon})$  time.*

Note that when combined with Observation 1, this implies a running time of  $\text{poly}(m, \frac{1}{\epsilon})$  to find an  $\epsilon$ -approximately stable lottery in the RANKING setting. Interestingly, assuming we can ask voters to compare any two committees, the algorithm for the APPROVAL SET and RANKING settings are exactly the same, and as the number of voters becomes large, the average number of queries asked to any individual voter goes to zero. Since the algorithm is the same in both cases, as a simple corollary, it also shows the existence of a stable solution when there is a mix of voters in the population, some with APPROVAL SET preferences and others with RANKING preferences.

We note that Theorem 1.2 does not follow by analyzing extant voting rules. For instance, in the APPROVAL SET setting, Aziz *et al.* Aziz et al. [2] showed that the voting rule, Proportional Approval Voting (PAV) [3, 25], that satisfies justified representation, fails to find a stable outcome. In PAV, if voter  $v$  approves  $r_v$  candidates in the committee, this voter is assigned a score  $s_v = 1 + \frac{1}{2} + \dots + \frac{1}{r_v}$ ; the winning committee maximizes  $\sum_v s_v$ . We strengthen their example to show that PAV cannot find better than an  $O(\sqrt{K})$ -approximately stable committee (Theorem 3.2 in Section 3). Further, as mentioned above, deterministic stability is simply not possible for the RANKING setting, which rules out trying to prove the above result via analyzing any deterministic voting rule. In that sense, we find the existence of stable lotteries quite surprising.

Though not the main focus of the paper, we finally consider the existence of deterministic stable committees in the APPROVAL SET setting. We make progress on this question and show that when the committee size is  $K \leq 3$ , a stable committee always exists (Theorem 3.1 in Section 3). Prior to this, the only results known (even for  $K = 3$ ) were the existence of committees that are approximately stable [10].

*Techniques.* The most interesting aspect of Theorem 1.2 is its proof, which uses the probabilistic method. We show that the question of whether stable lotteries exist reduces to deciding if a zero-sum game has negative value. The dual problem involves finding a stable solution given a lottery over blocking committees. We show its existence by developing a *rounding procedure* for the dual problem. This procedure performs *probability matching*, and simply chooses candidates with probability proportional to their marginal probability in the blocking lottery. We then argue that this rounding procedure always has negative expected value. Showing this for the APPROVAL SET setting requires proving a new deviation inequality (Lemma 1.3) for sums of random variables.

**LEMMA 1.3 (PROBABILITY MATCHING LEMMA).** *Let  $X$  be the sum of independent Bernoulli random variables:  $X = X_1 + X_2 + \dots + X_n$  where  $X_i \sim \text{Bernoulli}(1, p_i)$ . Let  $Y$  be any non-negative **integer** random variable with  $\mathbf{E}[Y] \leq \beta \mathbf{E}[X]$ . Then,*

$$\Pr[X < Y] < \beta .$$

Since we need each committee in the lottery to have size  $K$ , we combine this large deviation inequality with dependent rounding techniques [14] to preserve committee size. This rounding procedure can be plugged in as an oracle to the multiplicative weight update method [1], yielding the desired running time for the computational problem of finding an  $L$ -stable lottery. We believe

the template of our proofs via large deviation inequalities and probabilistic method may find further applications in the emerging theory of fair allocations.

*RoadMap.* We present the proof of Theorem 1.2 in Section 2. We split this proof into a generic rounding portion (Sections 2.1 and 2.2) and the subsequent part that is specific to APPROVAL SET (Section 2.3) and RANKING (Section 2.4). We present an efficient implementation in Section 2.5. In Section 3, we consider the notion of deterministic stability. We show the existence of stable committees for  $K = 3$ , and show that the PAV rule cannot be better than  $O(\sqrt{K})$ -approximately stable. We conclude with open questions in Section 4.

## 1.5 Related Work

We now compare our notion of stability with other notions of stability and fairness extant in literature. We first note that there is a rich literature on using lotteries to achieve fairness, including the influential papers of Hylland and Zeckhauser Hylland and Zeckhauser [15] and Bogomolnaia and Moulin Bogomolnaia and Moulin [4]. Our work is more in the spirit of the latter in the sense that our model does not require voters to specify cardinal utilities over committees.

*The Lindahl Equilibrium.* Our notion of stability is inspired by the notion of *core* in cooperative game theory and was first phrased in game theoretic terms by Scarf Scarf [24]. It has been extensively studied in public goods settings [9, 12, 20]. Much of this literature considers convex preferences, which translates to voters having preferences over lotteries, and deviating to another lottery if their expected utility increases. In other words, for additive utility functions  $u_v$ , we say a lottery  $\Delta'$  blocks  $\Delta$  if there is a coalition of voters of size at least  $n \cdot |\mathcal{S}'|/K$  such that for all voters  $v$  in this coalition, we have:

$$\mathbf{E}_{S' \sim \Delta'} [u_v(S')] \geq \mathbf{E}_{S \sim \Delta} [u_v(S)]$$

with at least one inequality strict. This notion builds on the seminal work of Foley Foley [12] on the Lindahl market equilibrium. In this equilibrium, each candidate is assigned a per-voter price. If the voters choose their utility maximizing allocation subject to spending a dollar, then (1) they all choose the same outcome; and (2) for each chosen candidate, the total money collected pays for that candidate. It can be shown via a fixed point argument that such an equilibrium pricing always exists when lotteries are allowed, and is a core outcome. Though this existence result is very general, holding for any compact convex set of preferences, it is not known how to compute such a core outcome efficiently even for the committee selection problem with ranking or approval set utilities.

Our notion of stability coincides with the notion of core if no randomization is allowed, but our notion of stable lotteries differs from how randomization is used in the core. In a Lindahl equilibrium, the voter compares the expected utility from the lottery with the utility on deviation, while in our notion, the lottery is first realized and subsequently the voters who see higher utility will deviate. In a sense, our notion justifies to any coalition of candidates that they do not have enough support given the current lottery, while the Lindahl equilibrium justifies to each coalition of voters that their utility is Pareto-optimal. This difference is subtle, but makes the two notions incomparable – the existence of one type of stable solution does not imply the other in any obvious way.

Our approach has two key advantages. First, our notion only requires voters to specify ordinal preferences over committees, and not over lotteries. Note that for both APPROVAL SET and RANKING settings, preferences over committees do not automatically imply preferences over lotteries – the latter must be explicitly specified via the choice of utility function. Secondly, our existence theorems, though specific to committee selection, have constructive proofs that lead to efficient algorithms

(while fixed point arguments in general do not). For instance, we can compute a stable lottery for RANKING in polynomial time, while we do not know how to compute a Lindahl equilibrium efficiently in this setting.

*Proportional Fairness.* A different notion of stability is the following: A deviating coalition of voters gets to choose a committee of size  $K$ , but its utility is scaled down proportional to the size of the coalition. In other words, given a coalition of voters of size  $\beta n$ , a committee  $S'$  of size  $K$  is blocking to a lottery  $\Delta$  if for all voters  $v$  in the coalition,

$$\beta u_v(S') \geq \mathbf{E}_{S \sim \Delta} [u_v(S)]$$

with at least one inequality strict. In this notion, it is relatively easy to show that maximizing the Nash product of voter utilities finds a stable lottery [9]; furthermore, the discrete analog of Nash welfare, proportional approval voting (PAV), finds an approximately stable solution for additive utilities [10]. Though this approach is computationally tractable, it explicitly needs voters to have cardinal utilities, and is otherwise incomparable to scaling down the committee size. Our notion is closer to how fairness has been thought about in the Economics literature on public goods, and to how it has been thought about in the multi-winner election literature. We have already delved into the former; we consider the latter next.

*Justified Representation.* Deterministic stability in the APPROVAL SET setting logically implies a number of fairness notions considered in multi-winner election literature with approval set preferences, such as justified representation, extended justified representation [2], and proportional justified representation [3, 23]. The idea behind all these proportional representation axioms is to define a notion of *cohesive* groups of agents that all approve a small set of candidates, and ensure that such groups of voters do not deviate, *i.e.*, that these groups are proportionally represented in the outcome. As mentioned above, justified representation is exactly the same as our 1-stability. Similarly, in extended justified representation (EJR), we only consider deviations by sub-populations of voters of size at least  $\alpha n$ , if they all approve the *same set* of  $\lfloor \alpha K \rfloor$  candidates. Therefore, a committee  $S$  satisfies EJR iff some voter in this sub-population approves at least  $\lfloor \alpha K \rfloor$  in  $S$ .

Unlike justified representation and its generalizations, stability is a general condition that holds for all coalitions of agents, not just those that are cohesive. We do pay a price for this generalization: our stability results are for lotteries of committees, and unlike justified representation and EJR, our algorithm for computing stable lotteries runs in polynomial time only when the deviating committee has constant size. As noted before, it is an open question whether (deterministic) stable committees exist in the APPROVAL SET setting.

*Committee Scoring Rules.* Elkind *et al.* Elkind *et al.* [8] show that given preferences of voters over candidates, a large class of voting rules for committee selection can be expressed as follows: For each voter, sort the positions of the committee members in its own ranking in non-decreasing order. Now apply a monotone scoring function to this vector of positions, and add up this score for all voters. The winning committee maximizes this score. This can be interpreted as assigning a utility function for each voter for each committee, and finding the committee that maximizes the sum of utilities over all voters, *i.e.*, the social optimum. In the RANKING setting, the voter derives utility only from the most preferred candidate in the committee; in fact, the Chamberlin-Courant voting rule Chamberlin and Courant [6] sums over all the voters, the Borda score of that voter's highest ranked candidate in the committee, and chooses the committee with the highest score. Other voting rules such as  $k$ -Borda assume more general submodular utilities, and it is an open question whether stable lotteries exist when voters' utilities are submodular.

Elkind *et al.* Elkind et al. [8] also show that for several natural scoring functions including Chamberlin-Courant, the resulting voting rule is NP-HARD. Further, since these voting rules are deterministic, none of them can produce a stable outcome. Our contribution is to show that for the RANKING setting, stable *lotteries* always exists and can be computed in polynomial time.

## 2 EXISTENCE OF STABLE LOTTERIES

In this section, we will prove Theorem 1.2. We will do this in a sequence of steps. First, we will formulate stability as a zero-sum game, and take its dual. Next, we will devise a rounding procedure to show the dual has negative value, which will imply the existence of a stable solution. Though the rounding procedure is the same for both APPROVAL SET and RANKING settings, we need different proofs of existence for each. In Section 2.5, we show that the rounding procedure also gives an efficient algorithm for finding an  $\epsilon$ -approximately stable lottery via the multiplicative weight method.

### 2.1 Dual Formulation of Stability

Recall the definition of a stable lottery from Definition 3. The existence of a stable lottery is equivalent to showing that:

$$\min_{\Delta} \max_{S'} \left( \mathbb{E}_{S \sim \Delta} [V(S, S')] - n \frac{|S'|}{K} \right) < 0 \quad (1)$$

where  $\Delta$  is a lottery over committees of size  $K$ .

We formulate a zero-sum game with an attacker and defender. The goal of the defender is to choose a stable committee, while the goal of the attacker is to find a blocking committee. The defender's pure strategies are committees  $S_d$  of size *exactly*  $K$ , and the attacker's pure strategies are committees  $S_a$  of size *at most*  $K$ . Denote the former set of strategies as  $\mathcal{S}_d$  and the latter by  $\mathcal{S}_a$ . Given a pair of pure strategies  $(S_d, S_a)$  the payoff to the attacker is:

$$Q(S_d, S_a) = V(S_d, S_a) - n \frac{|S_a|}{K} .$$

and the defender tries to minimize this value.

Therefore, the LHS of Equation (1) is the value of the game if the defender goes first and chooses mixed strategy  $\Delta_d$ , and the attacker subsequently chooses strategy  $S_a$  to maximize the value of the game. Therefore, the existence of a stable lottery is equivalent to asking whether the value of this zero-sum game is negative.

*Dual Formulation.* Suppose the attacker goes first and chooses a lottery  $\Delta_a$  over  $\mathcal{S}_a$ . Subsequently, the defender chooses a committee  $S_d \in \mathcal{S}_d$  to minimize the value of the game. Using duality in zero-sum games,

$$\min_{\Delta_d} \max_{S_a \in \mathcal{S}_a} \mathbb{E}_{S_d \sim \Delta_d} [Q(S_d, S_a)] = \max_{\Delta_a} \min_{S_d \in \mathcal{S}_d} \mathbb{E}_{S_a \sim \Delta_a} [Q(S_d, S_a)] .$$

Define  $\beta = \frac{\mathbb{E}_{S_a \sim \Delta_a} [|S_a|]}{K}$ . Note that  $\beta \in (0, 1]$  because  $\Delta_a$  is a lottery over  $\mathcal{S}_a$ , *i.e.*, over committees of size at most  $K$ . Then, we have:

$$\mathbb{E}_{S_a \sim \Delta_a} \left[ V(S_d, S_a) - n \frac{|S_a|}{K} \right] = \sum_{v \in \mathcal{N}} \left( \Pr_{S_a \sim \Delta_a} [S_a \succ_v S_d] \right) - \beta n = \sum_{v \in \mathcal{N}} \left( \Pr_{S_a \sim \Delta_a} [S_a \succ_v S_d] - \beta \right) .$$

Therefore, the value of the game is:

$$\max_{\Delta_a} \min_{S_d \in \mathcal{S}_d} \mathbb{E}_{S_a \sim \Delta_a} [Q(S_d, S_a)] = \max_{\Delta_a} \min_{S_d \in \mathcal{S}_d} \left( \sum_{v \in \mathcal{N}} \left( \Pr_{S_a \sim \Delta_a} [S_a \succ_v S_d] - \beta \right) \right) . \quad (2)$$



In order to prove the existence part of Theorem 1.2, we need to show that the RHS of the above identity is negative for the APPROVAL SET and RANKING settings.

## 2.2 Defender's Strategy: Probability Matching

Fix any mixed strategy  $\Delta_a$  for the attacker. Let  $\beta = \frac{E_{S_a \sim \Delta_a}[|S_a|]}{K}$ . We will construct a mixed defending strategy  $\Delta_d$  over  $\mathcal{S}_d$ , i.e. over committees of size  $K$ , so that

$$\forall v \in \mathcal{N} \quad \Pr_{S_a \sim \Delta_a, S_d \sim \Delta_d} [S_a \succ_v S_d] < \beta. \quad (3)$$

Equations (2) and (3) together show the existence of a defending strategy  $S_d$  that makes the attacker's payoff negative:

$$\min_{S_d \in \mathcal{S}_d} E_{S_a \sim \Delta_a} [Q(S_d, S_a)] \leq E_{S_d \sim \Delta_d} E_{S_a \sim \Delta_a} [Q(S_d, S_a)] = \sum_{v \in \mathcal{N}} \left( \Pr_{S_a \sim \Delta_a, S_d \sim \Delta_d} [S_a \succ_v S_d] - \beta \right) < 0.$$

Because such  $S_d$  exists for any mixed attacking strategy  $\Delta_a$ , the value of the game is negative, and hence a stable lottery always exists.

*Probability Matching.* We will construct the mixed defending strategy using *probability matching*. For candidate  $i \in C$ , let random variable  $Y_i = \mathbf{1}_{i \in S_a}$  be the indicator variable for the event  $i \in S_a$  when  $S_a \sim \Delta_a$ , and let  $p_i = E[Y_i] = \Pr_{S_a \sim \Delta_a} [i \in S_a]$ . Note that the inclusion of different candidates in the attacking strategy can be correlated, i.e., the  $Y_i$ 's can be correlated.

Let  $q_i = \min(1, p_i/\beta)$ . The defender's strategy will include  $i \in S_d$  with probability at least  $q_i$ , i.e., the defender will *probability match* the attacker for each candidate  $i \in C$ . We first make a simple observation.

CLAIM 1.  $\sum_{i \in C} q_i \leq K$ .

PROOF. The expected size of the attacker committee is  $\beta K$ , so  $\sum_i p_i = \beta K$ . Therefore,

$$\sum_i q_i = \sum_i \min(1, p_i/\beta) \leq \sum_i p_i/\beta = K. \quad \square$$

*Randomized Dependent Rounding.* We first increase  $\vec{q}$  arbitrarily to  $\vec{\alpha}$  so that  $\sum_i \alpha_i = K$  and  $\alpha_i \in [q_i, 1]$  for all  $i \in C$ . We will construct a mixed defending strategy  $\Delta_d$  over committees  $S_d$  of size  $K$  by performing *randomized dependent rounding* [14] on  $\vec{\alpha}$ . Let random variable  $X_i$  be the indicator variable for the event  $i \in S_d$  when  $S_d \sim \Delta_d$ . We omit the details of this rounding procedure, and instead summarize its properties in the following lemma:

LEMMA 2.1 (GANDHI ET AL. [14]). *The distribution  $\Delta_d$  obtained by performing randomized dependent rounding on  $\vec{\alpha}$  satisfies the following:*

- (1)  $\Delta_d$  is a distribution over committees of size  $K$ , i.e., for every  $S_d \sim \Delta_d$  we have  $|S_d| = K$ .
- (2)  $\Pr[X_i = 1] = \alpha_i$ , i.e., the marginal probability of choosing candidate  $i$  is preserved.
- (3) The random variables  $X_i$  are negatively correlated. More specifically, for every subset  $S \subseteq C$ ,

$$\Pr \left[ \bigwedge_{i \in S} X_i = b \right] \leq \prod_{i \in S} \Pr[X_i = b] \quad \forall b \in \{0, 1\}.$$

This completes the generic portion of the proof that there exists a stable lottery over committees. It remains to show that the defender's lottery  $\Delta_d$  we constructed satisfies Inequality (3). We will show this separately for the APPROVAL SET and RANKING settings.

### 2.3 Proof of Theorem 1.2 for APPROVAL SET Setting

We will now show Inequality (3) holds for the APPROVAL SET setting. As we show later, this statement is a stronger version of the Probability Matching Lemma (Lemma 1.3), where we replace independence of  $\{X_i\}$  by negative dependence.

Fix any voter  $v \in \mathcal{N}$  with approval set  $A_v$ . By the definition of APPROVAL SET:

$$\Pr_{S_a \sim \Delta_a, S_d \sim \Delta_d} [S_a \succ_v S_d] = \Pr_{S_a \sim \Delta_a, S_d \sim \Delta_d} [|S_a \cap A_v| > |S_d \cap A_v|] .$$

Suppose  $A_v = \{c_1, c_2, \dots, c_\ell, c_{\ell+1}, \dots, c_r\}$ . We simplify notation to denote  $\alpha_i = \alpha_{c_i}$ . These candidates are ordered so that  $\alpha_i$ 's are in ascending order. Further,  $\alpha_\ell < 1$  and  $\alpha_{\ell+1} = \alpha_{\ell+2} = \dots = \alpha_r = 1$ . Since  $c_{\ell+1}, \dots, c_r$  are selected in  $S_d$  with probability 1,

$$\Pr[|S_d \cap A_v| < |S_a \cap A_v|] \leq \Pr[|S_d \cap \{c_1, \dots, c_\ell\}| < |S_a \cap \{c_1, \dots, c_\ell\}|] .$$

Defining  $T_d := S_d \cap \{c_1, \dots, c_\ell\}$  and  $T_a := S_a \cap \{c_1, \dots, c_\ell\}$ , we have

$$\Pr_{S_a \sim \Delta_a, S_d \sim \Delta_d} [S_a \succ_v S_d] \leq \Pr[|T_d| < |T_a|] . \quad (4)$$

In order to show Inequality (3), our goal is therefore to show that  $\Pr[|T_d| < |T_a|] < \beta$ .

Recall that  $Y_i$  is the indicator random variable for whether  $c_i \in S_a$  when  $S_a \sim \Delta_a$ , and similarly,  $X_i$  is the indicator random variable for whether  $c_i \in S_d$  when  $S_d \sim \Delta_d$ . The  $Y_i$ 's can be correlated and  $E[Y_i] = p_i$ . For  $i \in \{1, 2, \dots, \ell\}$ , by Lemma 2.1, we have  $E[X_i] = \alpha_i \geq p_i/\beta$  and the  $X_i$ 's are negatively correlated. Define random variables  $X = \sum_{i=1}^{\ell} X_i$  and  $Y = \sum_{i=1}^{\ell} Y_i$ .

Using these notations, our goal is to prove

$$\Pr[|T_d| < |T_a|] = \Pr[X < Y] < \beta .$$

The next lemma uses Chernoff bound to upper bound the probability that  $X < \eta$ .

LEMMA 2.2. *Let  $X = \sum_{i=1}^{\ell} X_i$  and  $\mu = E[X]$ . For any integer  $\eta \geq 1$ ,  $\Pr[X < \eta] < \frac{\eta}{\mu}$ .*

PROOF. The statement holds trivially when  $\eta \geq \mu$ , so w.l.o.g. we can assume  $1 \leq \eta < \mu$ . Because  $\eta$  is an integer,  $\Pr[X < \eta] = \Pr[X \leq \eta - 1]$ .

By Lemma 2.1, the  $X_i$ 's produced by randomized dependent rounding [14] is negatively correlated, so we can apply Chernoff bound to the  $X_i$ 's [14, 21]:

$$\Pr[X \leq \eta - 1] = \Pr\left[X \leq \left(\frac{\eta - 1}{\mu}\right)\mu\right] \leq \exp\left(-\frac{1}{2}\left(1 - \frac{\eta - 1}{\mu}\right)^2 \mu\right) .$$

It remains to show that the RHS is less than  $\frac{\eta}{\mu}$  for any  $\mu > 0$  and any integer  $1 \leq \eta < \mu$ . One way to do so is to fix the ratio  $r = \frac{\eta}{\mu} \in (0, 1)$  and first minimize the RHS over  $\mu = \frac{\eta}{r} \geq \frac{1}{r}$ .

Let  $f(\mu) = (1 - r + \frac{1}{\mu})^2 \mu$ . We can compute its derivative:  $f'(\mu) = (1 - r)^2 - \frac{1}{\mu^2}$ . Thus,  $f(\mu)$  is decreasing when  $0 < \mu \leq \frac{1}{1-r}$  and increasing when  $\mu \geq \frac{1}{1-r}$ .

When  $r \geq \frac{1}{2}$ , we have  $\min_{\mu \geq \frac{1}{r}} f(\mu) = f(\frac{1}{1-r}) = 4(1 - r)$ , which implies  $\text{RHS} \leq \exp\left(-\frac{4(1-r)}{2}\right) < r$ . When  $r \leq \frac{1}{2}$ , we have  $\frac{1}{r} \geq \frac{1}{1-r}$ , so  $\min_{\mu \geq \frac{1}{r}} f(\mu) = f(\frac{1}{r}) = \frac{1}{r}$  and  $\text{RHS} \leq \exp\left(-\frac{1}{2r}\right) < r$ .  $\square$

Given Lemma 2.2, we can prove Inequality (3) by considering all possible values of  $Y$ :

$$\Pr[X < Y] = \sum_{\eta=1}^{\ell} \Pr[Y = \eta] \cdot \Pr[X < \eta] < \sum_{\eta=1}^{\ell} \Pr[Y = \eta] \frac{\eta}{E[X]} = \frac{E[Y]}{E[X]} \leq \beta .$$

The last step follows from the fact that  $E[X] = \sum_{i=1}^{\ell} \alpha_i \geq \sum_{i=1}^{\ell} p_i/\beta = E[Y]/\beta$ . This completes the proof that a stable lottery exists in the APPROVAL SET setting.

*Discussion.* The above proof also proves Lemma 1.3. Recall that Lemma 1.3 states that

$$\Pr[X < Y] < \beta$$

where  $X$  is the sum of independent Bernoulli random variables and  $Y$  is a non-negative integer random variable. Now since  $E[Y] \leq \beta E[X]$ , a simple application of Markov's inequality yields:

$$\Pr[E[X] < Y] < \beta .$$

However, this is not quite what we want. The quantity  $X = \sum_i X_i$  is the sum of independent Bernoulli random variables, and will deviate below its expectation. The proof in Section 2.3 varies the threshold  $\eta$ , and uses Chernoff bound to upper bound the probability that  $X$  deviates below  $\eta$ .

The probability matching approach chooses candidates in  $S_d$  (i.e., the variables  $X_i$ ) proportional to the marginals of the correlated distribution  $\Delta_a$  (i.e., the variables  $Y_i$ ). This is similar to the randomized contention resolution scheme for approximately maximizing submodular welfare [11]. However, this resemblance is superficial. First, we do not round with probability equal to the marginal, but instead need to scale it up by  $1/\beta$ . More importantly, the analysis for submodular welfare focuses on achieving a constant approximation against an LP bound, i.e., they show that when  $\beta = 1$ , then  $E[f(\vec{X})] \geq (1 - 1/e)E[f(\vec{Y})]$  for submodular function  $f$ . On the other hand, we show a more delicate statement  $\Pr[\sum_{i=1}^{\ell} Y_i > \sum_{i=1}^{\ell} X_i] < \beta$ . In fact, we can modify the tight example in [11] to show that for this specific rounding procedure, the statement  $\Pr[f(\vec{Y}) > f(\vec{X})] < \beta$  is *false* for arbitrary submodular functions  $f$ . This implies our probability matching procedure *cannot* be used to show the existence of stable lotteries when utilities of voters for committees are submodular. It is an open question to extend this proof to the case where the  $X_i \sim \text{Bernoulli}(s_i, \alpha_i)$ , which corresponds to the voters having arbitrary additive utilities for candidates.

## 2.4 Proof of Theorem 1.2 for RANKING Setting

We now prove Inequality (3) for the RANKING setting. Fix any voter  $v \in \mathcal{N}$ . Assume the candidates are ordered so that  $c_1 \succ_v c_2 \succ_v \dots \succ_v c_m$ . Let  $X_i$  denote the indicator random variable that the defender's committee  $S_d \sim \Delta_d$  includes  $c_i$ , and let  $Y_i$  denote the indicator random variable that the attacker's committee  $S_a \sim \Delta_a$  includes  $c_i$ .

Recall that the  $X_i$ 's are negatively correlated and  $Y_i$ 's can be arbitrarily correlated,  $E[Y_i] = p_i$ , and  $E[X_i] = \alpha_i \geq \min(1, p_i/\beta)$ . Let  $c_\ell$  denote the earliest (highest ranked) candidate in this ordering for whom  $\alpha_\ell = 1$ , i.e.,  $X_\ell = 1$  with probability 1. Note that for  $1 \leq j < \ell$ , we have  $\alpha_j < 1$  and thus  $\alpha_j \geq p_j/\beta$ .

Let  $c_j$  be the highest ranked candidate picked by the attacker. That is,  $j$  is a random variable which is the smallest index with  $Y_j = 1$ . Then,  $S_a \succ_v S_d$  iff  $X_1 = X_2 = \dots = X_j = 0$ , i.e., the defender fails to pick any candidate that ranks at least as high as  $j$ .

Because the defender will always pick  $c_\ell$ , for  $S_a \succ_v S_d$  to happen, the attacker's highest ranked candidate has to appear before  $c_\ell$  (i.e.,  $j < \ell$ ). Therefore,

$$\begin{aligned} \Pr[S_a \succ_v S_d] &= \sum_{j=1}^{\ell-1} \Pr[Y_1 = Y_2 = \dots = Y_{j-1} = 0 \text{ and } Y_j = 1 \text{ and } X_1 = X_2 = \dots = X_j = 0] \\ &\leq \sum_{j=1}^{\ell-1} \Pr[Y_j = 1 \text{ and } X_1 = X_2 = \dots = X_j = 0] . \end{aligned}$$

Since  $\{X_i\}$  is independent of  $\{Y_i\}$ , we have:

$$\begin{aligned} \Pr[S_a \succ_v S_d] &\leq \sum_{j=1}^{\ell-1} \Pr[Y_j = 1] \cdot \Pr[X_1 = X_2 = \dots = X_j = 0] \\ &\leq \sum_{j=1}^{\ell-1} \Pr[Y_j = 1] \cdot \Pr[X_1 = 0] \Pr[X_2 = 0] \dots \Pr[X_j = 0] \\ &\leq \sum_{j=1}^{\ell-1} \beta \alpha_j \prod_{i=1}^j (1 - \alpha_i). \end{aligned}$$

Here, the second inequality follows from the negative correlation of  $\{X_i\}$ . More precisely, this inequality is a special case of the negative correlation condition of Lemma 2.1. The final inequality follows since  $\mathbb{E}[Y_i] = p_i \leq \beta \alpha_i$ , and  $\mathbb{E}[X_i] = \alpha_i$ .

Consider the following stopping process: Consider candidates in order  $c_1, c_2, \dots, c_\ell$ . We stop at candidate  $c_1$  with probability  $\alpha_1$ ; if we do not stop, we stop at  $c_2$  independently with probability  $\alpha_2$ , and so on. Because  $\alpha_j < 1$  for all  $1 \leq j < \ell$ , we stop with probability less than 1 before we reach  $c_\ell$ :

$$\sum_{j=1}^{\ell-1} \alpha_j \prod_{i=1}^{j-1} (1 - \alpha_i) < 1.$$

Comparing this inequality with the previously obtained bound, we have:

$$\Pr[S_a \succ_v S_d] \leq \sum_{j=1}^{\ell-1} \beta \alpha_j \prod_{i=1}^j (1 - \alpha_i) \leq \beta \sum_{j=1}^{\ell-1} \alpha_j \prod_{i=1}^{j-1} (1 - \alpha_i) < \beta.$$

This completes the proof of Inequality (3) for the RANKING setting.

## 2.5 Computing Stable Lotteries Efficiently

In this section, we turn to the problem of computing a stable lottery efficiently, and prove the algorithmic part of Theorem 1.2. We will show a running time of  $\text{poly}(m^L, 1/\epsilon)$  to compute an  $\epsilon$ -approximately  $L$ -stable lottery (see Definition 3). This yields a polynomial running time for  $L$ -stable lotteries for constant  $L$  in the APPROVAL SET setting, and for stable lotteries in the RANKING setting. Since our algorithm is *the same* for both APPROVAL SET and RANKING settings (assuming voters have the ability to compare two committees); it also yields a stable lottery when there is a mix of voters, some with APPROVAL SET and others with RANKING preferences.

First observe that since the size of the strategy sets  $\mathcal{S}_d$  and  $\mathcal{S}_a$  are  $O(m^K)$ , and since all we are doing is solving a zero-sum game, there is clearly a  $\text{poly}(n, m^K)$  time algorithm. The key observation is that the randomized dependent rounding procedure allows us to compute a defender strategy efficiently, and this allows us to compute  $L$ -stable lotteries in time that depends on  $m^L$  instead of  $m^K$ . We then combine this with estimating the  $V(S, S')$  values by sampling voters and asking them to compare  $S$  and  $S'$ , thereby eliminating the dependence of the running time on  $n$ . Since the details follow from fairly standard ideas, we only outline the argument.

Given a lottery  $\Delta_a$  over  $\mathcal{S}_a$ , let  $\text{ORACLE}(\Delta_a, \epsilon)$  be a procedure that finds  $S_d \in \mathcal{S}_d$  such that

$$R_\epsilon(S_d, \Delta_a) \equiv \mathbb{E}_{S_a \sim \Delta_a} \left[ V(S_d, S_a) - (1 + \epsilon)n \frac{|S_a|}{K} \right] < 0.$$

**CLAIM 2.** For  $\epsilon \in (0, 1/5)$ ,  $\text{ORACLE}(\Delta_a, \epsilon)$  can be implemented in expected time  $\text{poly}(m^L, \frac{1}{\epsilon})$ , and is correct with high probability.

PROOF. First note that given  $\Delta_a$ , the randomized dependent rounding procedure in [14] can be implemented in  $O(m^2)$  time. The randomized dependent rounding outputs a distribution  $\Delta_d$  that satisfies  $\mathbb{E}_{T_d \sim \Delta_d} [R_0(T_d, \Delta_a)] < 0$ . We have:

$$\mathbb{E}_{T_d \sim \Delta_d} [\mathbb{E}_{S_a \sim \Delta_a} [V(T_d, S_a)]] \leq \mathbb{E}_{S_a \sim \Delta_a} \left[ n \frac{|S_a|}{K} \right].$$

By Markov's inequality, we have:

$$\Pr_{T_d \sim \Delta_d} \left[ \mathbb{E}_{S_a \sim \Delta_a} [V(T_d, S_a)] \geq (1 + \epsilon) \mathbb{E}_{S_a \sim \Delta_a} \left[ n \frac{|S_a|}{K} \right] \right] < \frac{1}{1 + \epsilon}. \quad (5)$$

Given an output  $T_d$  of randomized dependent rounding, for each  $S_a$  of size at most  $L$ , we can estimate  $V(T_d, S_a)$  to within an additive  $\frac{\epsilon n}{K}$  as  $\hat{V}(T_d, S_a)$ , by sampling  $\text{poly}(m, \frac{1}{\epsilon})$  voters and asking them to compare  $T_d$  with  $S_a$ . This allows us to estimate  $\mathbb{E}_{S_a \sim \Delta_a} [V(T_d, S_a)]$  to within an additive  $\frac{\epsilon n}{K}$  with high probability in  $\text{poly}(m^L, \frac{1}{\epsilon})$  time, since the support of  $\Delta_a$  has size  $O(m^L)$ . Since the entire algorithm takes  $\text{poly}(m, 1/\epsilon)$  steps, we can therefore assume that with probability at least  $1 - 1/m$ , our estimate of any  $V(T_d, S_a)$  is accurate for all steps of the algorithm. Call the resulting estimate of  $R_\epsilon(T_d, \Delta_a)$ , where each  $V(T_d, S_a)$  is estimated by sampling, as  $\hat{R}_\epsilon(T_d, \Delta_a)$ .

We first show that for  $T_d \sim \Delta_d$ , the event  $\hat{R}_{2\epsilon}(T_d, \Delta_a) < 0$  happens with probability at least  $\epsilon/2$ . Inequality (5) implies  $R_\epsilon(T_d, \Delta_a) < 0$  with probability at least  $\frac{\epsilon}{1+\epsilon}$ . When this event happens, we have  $R_{2\epsilon}(T_d, \Delta_a) \leq R_\epsilon(T_d, \Delta_a) - \frac{\epsilon n}{K} < -\frac{\epsilon n}{K}$ . Since our sampled estimates of  $\mathbb{E}_{S_a \sim \Delta_a} [V(T_d, S_a)]$  are accurate within an additive  $\frac{\epsilon n}{K}$  with high probability, we have  $\hat{R}_{2\epsilon}(T_d, \Delta_a) \leq R_{2\epsilon}(T_d, \Delta_a) + \frac{\epsilon n}{K} < 0$ .

Therefore, if the randomized dependent rounding procedure is repeated till  $\hat{R}_{2\epsilon}(T_d, \Delta_a) < 0$ , it takes  $\text{poly}(\frac{1}{\epsilon})$  trials in expectation. We conclude the proof by noting that the resulting  $T_d$  is a feasible solution for  $\text{ORACLE}(\Delta_a, 3\epsilon)$  with high probability:  $R_{3\epsilon}(T_d, \Delta_a) \leq R_{2\epsilon}(T_d, \Delta_a) - \frac{\epsilon n}{K} \leq \hat{R}_{2\epsilon}(T_d, \Delta_a) < 0$ .  $\square$

We now use the multiplicative weight update (MWU) method [1] in a standard fashion:

- (1) Given the adversary's strategy  $\Delta_a^t$  at time  $t$ , run  $\text{ORACLE}(\Delta_a^t, \epsilon)$  to compute  $S_d^t$  satisfying  $R_\epsilon(S_d^t, \Delta_a^t) < 0$ .
- (2) Treat each committee  $S_a$  of size at most  $L$  as an expert, and set its gain to be  $g^t(S_a) = V(S_d^t, S_a) - (1 + \epsilon)n \frac{|S_a|}{K}$ . Again, these gains can be approximately computed by sampling the voters and asking them to compare  $S_d^t$  and  $S_a$ .
- (3) Feed these gains to the MWU algorithm, which outputs a lottery  $\Delta_a^{t+1}$  over the experts.

By a standard analysis, if we run the procedure for  $T = \text{poly}(m, \frac{1}{\epsilon})$  steps, the lottery  $\Delta_d$  that chooses  $S_d^t$  for  $t = 1, 2, \dots, T$  with equal probability will be an  $\epsilon'$ -approximately  $L$ -stable lottery for  $\epsilon' = O(\epsilon)$ . Further, the algorithm only involves asking voters to compare two committees, and is hence the same for both APPROVAL SET and RANKING.

### 3 DETERMINISTIC STABILITY FOR APPROVAL SET SETTING

So far, we have considered existence of stable lotteries. As mentioned earlier, for the APPROVAL SET setting, the existence of stable (deterministic) committees is a tantalizing open question. In this section, we present some results that make progress towards the goal of finding stable committees. First, we show that stable committees always exist when the committee size is  $K \leq 3$ , regardless of the number of candidates and voters. In addition, we show that the PAV rule that satisfies justified representation, fails to find *any non-trivial* approximation to stable committees. This strengthens the negative result in [2] to include inapproximability.

### 3.1 Existence of Stable Committees When $K = 3$

The next theorem states that for APPROVAL SET setting and  $K = 3$ , stable committee always exists and can be computed efficiently. Recall that  $m$  is the number of candidates,  $n$  is the number of voters, and  $V(S, S')$  is the number of voters who prefer committee  $S'$  over  $S$ .

**THEOREM 3.1.** *In the APPROVAL SET setting, for committee size  $K = 3$ , a stable committee always exists. Moreover, it can be computed in time  $O(m^3n)$ .*

**PROOF.** For a committee  $S$ , we use  $n_i(S)$  for the number of voters who approve exactly  $i$  candidates in  $S$ , and we use  $n_{\geq i}(S)$  for the number of voters who approve at least  $i$  candidates in  $S$ .

First observe that it is sufficient to find committee  $T$  with no blocking committees of size 1 or 2. This is because either  $T$  is stable, or it has blocking committees of size 3 that Pareto-dominate  $T$ . Consider any Pareto-optimal committee  $T'$  that Pareto-dominates  $T$ .  $T'$  is stable because it has no blocking committees of size 3 due to its Pareto-optimality; and since  $T'$  makes all voters happier than they were under  $T$ ,  $T'$  has no blocking committees of size 1 or 2 either.

The rest of the proof shows that such committee  $T$  always exists via the following case analysis:

- (1) There exists a committee  $S = \{a, b\}$  such that  $n_2(S) > \frac{n}{3}$ .
- (2) There are no committees of size 2 that satisfy (1), but there exists a committee  $S = \{a, b\}$  such that  $n_{\geq 1}(S) \geq \frac{2n}{3}$ .
- (3) There are no committees of size 2 that satisfy (1) or (2).

For Case (1), assume w.l.o.g. that  $S$  has a blocking committee of size 1 or 2. (Otherwise we can choose  $T = S$ .) However, any committee  $S'$  of size 2 is not blocking, because  $v(S, S') \leq n - n_2(S) < \frac{2n}{3}$ . Therefore, it must be the case that there is a blocking committee  $S' = \{c\}$  of size one. We add  $c$  to  $S$  and consider the committee  $T = \{a, b, c\}$ . Note that  $n_{\geq 2}(T) \geq n_2(S) > \frac{n}{3}$ , and  $n_1(T) \geq v(S, \{c\}) \geq \frac{n}{3}$  because voters who prefer  $\{c\}$  over  $S$  must approve  $c$  but not anyone in  $S = \{a, b\}$ . We argue that  $T$  satisfies our requirements, because for any committee  $T'$  of size 2, we have  $V(T, T') \leq n - n_{\geq 2}(T) < \frac{2n}{3}$ ; and for any committee  $T'$  of size 1, we have  $V(T, T') \leq n_0(T) = n - n_{\geq 2}(T) - n_1(T) < \frac{n}{3}$ .

For Case (2), assume w.l.o.g. there is a candidate  $c$  with  $v(S, \{c\}) > 0$ . Consider the committee  $T = \{a, b, c\}$ . Note that  $n_{\geq 1}(T) = n_{\geq 1}(S) + V(S, \{c\}) > \frac{2n}{3}$ . We argue that  $T$  satisfies our requirements, because for any  $T'$  of size 1, we have  $v(T, T') \leq n_0(T) = n - n_{\geq 1}(T) < \frac{n}{3}$ ; and for any  $T'$  of size 2, we have  $V(T, T') \leq n_2(T') + n_0(T) < \frac{2n}{3}$ . The last inequality is because voters who prefer  $T'$  over  $T$  must either approve 2 candidates in  $T'$ , or approve 0 candidates in  $T$ ; and  $n_2(T') \leq \frac{n}{3}$  since we are not in Case (1).

Finally, in Case (3), there are no blocking committee  $T'$  of size 2 because  $V(\emptyset, T') = n_{\geq 1}(T') < \frac{2n}{3}$ . This allows us to focus only on blocking committees of size 1. There is w.l.o.g. some candidate  $a$  such that  $V(\emptyset, \{a\}) \geq \frac{n}{3}$ , otherwise we can set  $T = \emptyset$ . Then again, there is w.l.o.g. some candidate  $b$  such that  $V(\{a\}, \{b\}) \geq \frac{n}{3}$ , otherwise we can set  $T = \{a\}$ . However, this contradicts the assumption that we are in Case (3), because  $n_{\geq 1}(\{a, b\}) = V(\emptyset, \{a\}) + V(\{a\}, \{b\}) \geq \frac{2n}{3}$ .

It takes time  $O(m^2n)$  to find a committee  $T$  with no size 1 and 2 blocking committees. This can be done by enumerating all committees of size 2 in time  $O(m^2n)$  to decide which case we are in, and then the bottleneck is Case (1) in which we can find a blocking committee of size 1 in time  $O(mn)$ . Starting from  $T$ , we can find a stable solution  $T'$  in time  $O(m^3n)$  by enumerating all committees of size 3 and maintaining the current Pareto-optimal solution. Hence, the overall running time is  $O(m^2n + mn + m^3n) = O(m^3n)$ .  $\square$

### 3.2 Proportional Approval Voting Can Be Far From Stable

Recall that the PAV rule works as follows: For any committee  $S$ , suppose voter  $v$  approves  $r$  candidates in  $S$ , then we set the score  $q_v(S) = 1 + \frac{1}{2} + \dots + \frac{1}{r}$ . The PAV rule finds  $S$  that maximizes  $\sum_v q_v(S)$ . We show that the PAV rule cannot be better than  $O(\sqrt{K})$ -approximately stable.

**THEOREM 3.2.** *In the APPROVAL SET setting, the PAV rule may output a committee that is not  $o(\sqrt{K})$ -approximately stable.*

**PROOF.** Consider the following example (illustrated in Figure 1 where each rectangle is a candidate whose projection on  $x$ -axis corresponds to the voters who approve her): The set of voters is divided into two equal-sized disjoint sets  $\mathcal{N}_L$  and  $\mathcal{N}_R$ . There are 4 different sets of candidates:

- (1) Set A: There are  $\frac{P}{2}$  ( $P$  is a parameter related to  $K$  to be set later) candidates in A. The voters who approve a candidate in A are exactly those voters in  $\mathcal{N}_L$ .
- (2) Set B: There are  $\frac{P}{2}$  candidates in B. The voters who approve a candidate in B are exactly those voters in  $\mathcal{N}_R$ .
- (3) Set C: There are  $\frac{P}{2} \cdot \frac{P}{4}$  candidates in C. Voters in  $\mathcal{N}_R$  do not approve candidates in C. Every candidate in C is approved by  $\frac{N_L}{P/4}$  voters. Each voter in  $\mathcal{N}_L$  approves  $\frac{P}{2}$  candidates in C.
- (4) Set D: There are  $\frac{P}{2}$  candidates in D. Voters in  $\mathcal{N}_L$  do not approve candidates in D. Every candidate in D is approved by  $\frac{N_R}{P/2}$  voters. Each voter in  $\mathcal{N}_R$  approves one candidate in D.

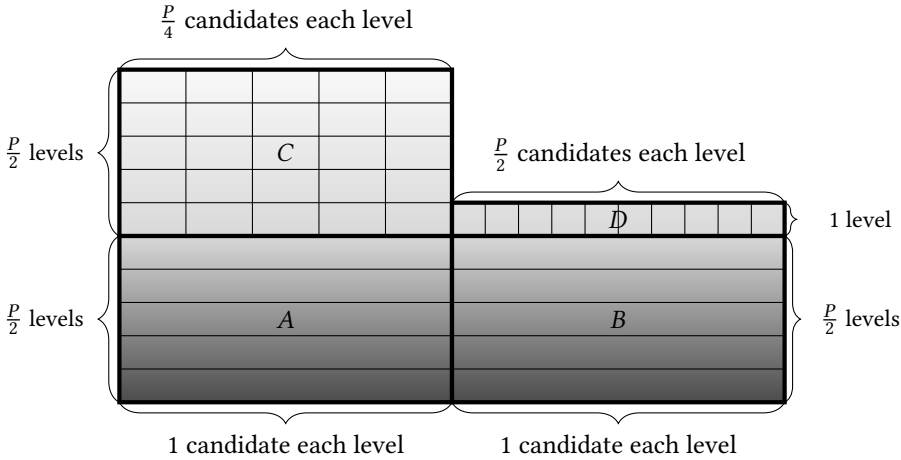


Fig. 1. PAV cannot be better than  $O(\sqrt{K})$ -approximately stable.

Consider running the PAV rule on this instance when  $K = P + \frac{P^2}{8}$ . PAV will first choose all candidates in A and B. Note that at this point, any candidate in C has marginal contribution at least  $\frac{1}{P} \cdot \frac{N_L}{P/4}$ , while any candidate in D has marginal contribution at most  $\frac{1}{(P/2)+1} \cdot \frac{N_R}{P/2}$ , which is strictly smaller. Therefore, PAV will select the committee  $S = A \cup B \cup C$ .

However, voters in  $\mathcal{N}_R$  can form a coalition and deviate to committee  $S' = B \cup D$ . All voters in  $\mathcal{N}_R$  are better off, thus  $V(S, S') = |\mathcal{N}_R| = \frac{n}{2}$ . For PAV to be  $\epsilon$ -approximately stable, we need

$$\frac{n}{2} \leq (1 + \epsilon) \frac{|S'|}{K} \cdot n = (1 + \epsilon) \frac{P}{K} \cdot n.$$

Since  $K = P + \frac{P^2}{8}$ , we need  $\epsilon = \Omega(\sqrt{K})$ . □

## 4 CONCLUSIONS

We view our results as a first step in understanding the general notion of stability for committee selection problems. We conclude with several open questions. The first major open question is the existence of deterministic stable committees in the APPROVAL SET setting, generalizing our positive result for  $K = 3$  to general  $K$ . We conjecture that such a stable committee always exists. Via computer-assisted search, we have shown that this conjecture holds for small numbers of voters and candidates ( $m + n \leq 14$ ). However, as we have seen, existing voting rules seem incapable of finding such a committee, which makes this question very tantalizing.

The next open question is whether a stable lottery exists for ordinal preferences over committees that result from more general cardinal preferences. Our “holy grail” is to have *one algorithm* that finds a stable lottery using ordinal preferences over committees, that works for a wide range of underlying cardinal utilities. One immediate extension of our work will be to the case where the utility of a voter is *additive* in the set of candidates in the committee. An obvious approach for attacking this problem is to extend Lemma 1.3 to the case where  $Y_i \sim \text{Bernoulli}(s_i, p_i)$ , but this will require new ideas.

Finally, our algorithm for  $L$ -stable lotteries has running time that depends on  $m^L$ . Though this is not an issue for RANKING, it would be good to remove this dependence for APPROVAL SET and develop a poly-time algorithm for any  $L$ , or show that this is not possible.

## ACKNOWLEDGMENTS

Kamesh Munagala is supported by NSF grants CCF-1408784, CCF-1637397, and IIS-1447554; and research awards from Adobe and Facebook. Kangning Wang is supported by NSF grants IIS-1447554 and CCF-1637397. This work was done while Zhihao Jiang was visiting Duke University.

## REFERENCES

- [1] Sanjeev Arora, Elad Hazan, and Satyen Kale. 2012. The Multiplicative Weights Update Method: a Meta-Algorithm and Applications. *Theory of Computing* 8, 6 (2012), 121–164.
- [2] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- [3] H. Aziz, E. Elkind, S. Huang, M. Lackner, L. Sanchez Fernandez, and P. Skowron. 2018. On the Complexity of Extended and Proportional Justified Representation. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [4] Anna Bogomolnaia and Hervé Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100, 2 (2001), 295 – 328.
- [5] Steven J. Brams, D. Marc Kilgour, and M. Remzi Sanver. 2007. A Minimax Procedure for Electing Committees. *Public Choice* 132, 3 (2007), 401–420.
- [6] John R. Chamberlin and Paul N. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *The American Political Science Review* 77, 3 (1983), 718–733.
- [7] H. R. Droop. 1881. On Methods of Electing Representatives. *Journal of the Statistical Society of London* 44, 2 (1881), 141–202.
- [8] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. 2017. Properties of Multiwinner Voting Rules. *Social Choice and Welfare* 48, 3 (2017), 599–632.
- [9] B. Fain, A. Goel, and K. Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*. 384–399.
- [10] Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*. 575–592.
- [11] Uriel Feige and Jan Vondrák. 2010. The Submodular Welfare Problem with Demand Queries. *Theory of Computing* 6, 11 (2010), 247–290.
- [12] D. K. Foley. 1970. Lindahl’s Solution and the Core of an Economy with Public Goods. *Econometrica* 38, 1 (1970), 66–72.
- [13] Eric Friedman, Vasilis Gkatzelis, Christos Alexandros Psomas, and Scott Shenker. 2019. Fair and Efficient Memory Sharing: Confronting Free Riders. In *33<sup>rd</sup> AAAI Conference on Artificial Intelligence (AAAI)*.
- [14] R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. 2006. Dependent Rounding and Its Applications to Approximation Algorithms. *J. ACM* 53, 3 (2006), 324–360.



- [15] Aanund Hylland and Richard Zeckhauser. 1979. The Efficient Allocation of Individuals to Positions. *Journal of Political Economy* 87, 2 (1979), 293–314.
- [16] M. Kunjir, B. Fain, K. Munagala, and S. Babu. 2017. ROBUS: Fair Cache Allocation for Data-parallel Workloads. In *Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD)*. 219–234.
- [17] Tyler Lu and Craig Boutilier. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI) (IJCAI'11)*. 280–286.
- [18] Reshef Meir, Ariel D. Procaccia, Jeffrey S. Rosenschein, and Aviv Zohar. 2008. Complexity of Strategic Behavior in Multi-winner Elections. *J. Artif. Int. Res.* 33, 1 (Sept. 2008), 149–178.
- [19] Burt L. Monroe. 1995. Fully Proportional Representation. *The American Political Science Review* 89, 4 (1995), 925–940.
- [20] T. J Muench. 1972. The Core and the Lindahl Equilibrium of an Economy with a Public Good: An Example. *Journal of Economic Theory* 4, 2 (1972), 241 – 255.
- [21] Alessandro Panconesi and Aravind Srinivasan. 1997. Randomized Distributed Edge Coloring via an Extension of the Chernoff-Hoeffding Bounds. *SIAM J. Comput.* 26, 2 (1997), 350–368.
- [22] Ariel D. Procaccia, Jeffrey S. Rosenschein, and Aviv Zohar. 2008. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare* 30, 3 (01 Apr 2008), 353–362.
- [23] L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. 2017. Proportional Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 670–676.
- [24] H. E. Scarf. 1967. The Core of an N Person Game. *Econometrica* 35, 1 (1967), pp. 50–69.
- [25] T. N. Thiele. 1895. Om flerfoldsvælg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*. 415–441.