

Multiwinner Voting with Fairness Constraints

L. Elisa Celis, Lingxiao Huang, Nisheeth K. Vishnoi

EPFL, Switzerland

first.last@epfl.ch

Abstract

Multiwinner voting rules are used to select a small representative subset of candidates or items from a larger set given the preferences of voters. However, if candidates have sensitive attributes such as gender or ethnicity (when selecting a committee), or specified types such as political leaning (when selecting a subset of news items), an algorithm that chooses a subset by optimizing a multiwinner voting rule may be unbalanced in its selection – it may under or over represent a particular gender or political orientation in the examples above. We introduce an algorithmic framework for multiwinner voting problems when there is an additional requirement that the selected subset should be “fair” with respect to a given set of attributes. Our framework provides the flexibility to (1) specify fairness with respect to *multiple, non-disjoint* attributes (e.g., ethnicity *and* gender) and (2) specify a score function. We study the computational complexity of this constrained multiwinner voting problem for monotone and submodular score functions and present several approximation algorithms and matching hardness of approximation results for various attribute group structure and types of score functions. We also present simulations that suggest that adding fairness constraints may not affect the scores significantly when compared to the unconstrained case.

1 Introduction

The problem of selecting a committee from a set of candidates given the preferences of voters is called multiwinner voting and arises in various social, political, and e-commerce settings; from electing a parliament, to choosing a committee, to selecting products to display. Formally, there is a set C of m “candidates” that can be selected (i.e., people, products, or other items) and a set A of n “voters” who have (possibly incomplete) preference list over the m candidates; we assume these lists are given. The goal is to select a subset of C of size k based on these preferences.

It remains to specify how the selection will be made. One common approach is to define a “total” score function $\text{score}_A : 2^C \rightarrow \mathbb{R}_{\geq 0}$ which gives a score that depends on the voters’ preferences to each potential committee. If the voter set A is clear from the context, we use score for short. This reduces the selection to an optimization problem: choose a

committee of size k that maximizes the score. Different views on the desired properties of the selection process have led to a number of different scoring rules and, consequently, to a variety of different algorithmic problems. Prevalent examples include committee scoring rules [Aziz *et al.*, 2017b; Elkind *et al.*, 2014], approval-based rules [Aziz *et al.*, 2017a], OWA-based rules [Skowron *et al.*, 2015b], variants of the Monroe rule [Betzler *et al.*, 2013; Monroe, 1995; Skowron *et al.*, 2015a] and the goalbase rules [Uckelman *et al.*, 2009].

However, it has been shown that voting rules, in the most general sense, can create or propagate biases; e.g., negatively affecting the fraction of women in the US legislature [Representation2020, 2017], and resulting in an electorate that under-represents minorities [Faliszewski *et al.*, 2016]. Furthermore, such algorithmic biases have been shown to influence and reinforce human stereotypes [Kay *et al.*, 2015]. An increasing awareness of such problems has led governments to generic [Bundy, 2017] and specific [Zealand, 1986] recommendations that aim to ensure sufficient representation of minority populations.

In response, “proportional representation” rules [Monroe, 1995], that represent the electorate proportionately in the elected body, have been developed. Formally, let $P_1, \dots, P_p \subseteq [m]$ be p *disjoint* groups of candidates where $i \in \{1, \dots, p\}$ represents a given group. For any $i \in [p]$, let f_i be the fraction of voters who belong to (or, more generally, prefer) group i . Then, a voting rule achieving full proportionality would ensure that the selected committee S satisfies $[k \cdot f_i] \leq |S \cap P_i| \leq \lceil k \cdot f_i \rceil$; [Brill *et al.*, Definition 5]. Proportional representation schemes include the proportional approval voting rule [Brill *et al.*,] and the Chamberlain Courant rule [1983]. Other approaches consider different notions of “fairness” in representation; e.g., in degressive proportionality it is argued that minorities should have disproportionately many representatives.

We present the first algorithmic framework for multiwinner voting problems that can (1) incorporate general notions of fairness with respect to arbitrary group structures (e.g., satisfying fairness across multiple attributes simultaneously) and (2) outputs a subset that maximizes the given score function subject to these constraints. This requires the development of new approximate algorithms (Section 6) and hardness of approximation results (Section 5) which depend on both the score function and the structure of the fairness constraints. Empirically, we show that existing multiwinner voting rules may introduce bias and that our approach not only ensures fairness, but does so with a score that is close to the (unconstrained) optimal (see Section 7).

2 Our Contributions

Model. In a multiwinner voting setting, we are given m candidates, n voters that each has a (potentially incomplete and/or non-strict) preference list over the m candidates, a score function $\text{score} : 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$ defined by these lists, and a desired number $k \in [m]$ of winners. In addition, to consider fair solutions, we are given arbitrary (potentially non-disjoint) groups of candidates $P_1, \dots, P_p \subseteq C$, and fairness constraints on the selected winner set S of the form: $\ell_i \leq |S \cap P_i| \leq u_i, \forall i \in [p]$ for given numbers $\ell_1, \dots, \ell_p, u_1, \dots, u_p \in \mathbb{Z}_{\geq 0}$.

The goal of the constrained multiwinner voting problem is to select a committee of size k that maximizes $\text{score}(S)$ and satisfies all fairness constraints. If the score function is monotone and submodular we call the problem the *constrained MS multiwinner voting problem*. This includes the well-studied Chamberlin–Courant (CC) rule, the Monroe rule, the OWA-based rules and the goalbase rules.

Results. The algorithmic problems that arise largely remain NP-hard, hence we focus on developing approximation algorithms. An important practical parameter, that also plays a role in the complexity of the constrained multiwinner voting problem, is the maximum number of groups to which a candidate can belong; we denote it by Δ . In real-world situations, we expect Δ to be a small constant, i.e., each candidate only belongs to a few groups. Our main results (classified by the kind of fairness parameters) are summarized in Tables 1 and 2.

When $\Delta = 1$, e.g., when the groups partition the candidates, we present a $(1 - 1/e)$ -approximation algorithm (Theorem 14) which is optimal given the $(1 - 1/e - \varepsilon)$ -hardness of approximation result by [Nemhauser *et al.*, 1978]. However, when $\Delta \geq 3$, unlike the unconstrained case, even checking whether there is a feasible solution becomes NP-hard (Theorem 6). The problem of finding a solution that violates cardinality or fairness constraints up to any multiplicative constant factor remains NP-hard (Theorem 7 and 8). Moreover, even if the feasibility is guaranteed, the problem remains hard to approximate within a factor of $\Omega(\log \Delta / \Delta)$ (Theorem 9).

To bypass this issue, we assume that the problem instance always has a feasible solution, and suggest natural sufficient conditions that *guarantee* feasibility. For instance, if the fraction of each group in the selected committee is allowed some slack as compared to their proportion in the set of candidates ($\ell_i \leq k(|P_i|/m - 0.05)$ and $u_i \geq k(|P_i|/m + 0.05)$), then a random committee of size k is feasible with high probability. When feasibility is guaranteed, for the class of MS voting rules we give a near-optimal bi-criterion approximation algorithm that only violates each fairness constraint by a small multiplicative factor (Theorem 11).

Techniques The algorithmic results combine two existing tools that have been extensively used in the monotone submodular maximization literature. The first, “multilinear extension” (Definition 10), extends the discrete MS score function to a continuous function over a relaxed domain. By applying a continuous greedy process via multilinear extension, a fractional solution with a high score can be computed efficiently. The second is to round the fractional solution to a

committee of size k by “dependent rounding”. In the case of $\Delta \geq 2$ (Theorem 11), we use a swap randomized rounding procedure introduced by [Chekuri *et al.*, 2010]. In the case of $\Delta = 1$ (Theorem 14) we design a two-layered dependent rounding procedure that runs in linear time. Some of the algorithmic results are achieved by reduction to well-studied problems, like the monotone submodular maximization problem with Δ -extendible system (Theorem 12) and constrained set multi-cover (Theorem 13). The hardness results follow from reductions from well-known NP-hard problems, including Δ -hypergraph matching, 3-regular vertex cover, constrained set multi-cover and independent set, and borrow techniques from a recent work on fairness for ranking problems [Celis *et al.*, 2018b].

We also study special cases; setting in which the fairness constraints involve only lower bounds (Theorem 13) or only upper bounds (Theorem 12). We further consider a certain specific MS score functions such as SNTV [Cox, 1994], or α -CC and β -CC [Chamberlin and Courant, 1983] where the unconstrained problem has recently received considerable attention. See Table 2 for a summary of these results, and natural corollaries which are described in the full version of the paper [Celis *et al.*, 2017b].

Generality. This approach is general in that 1) it can handle arbitrary MS score functions, 2) multiple sensitive attributes which can take on arbitrary group structures, 3) interval constraints that need not specify exact probabilities of representations for each group, and in doing so 4) can satisfy many different existing notions of fairness.

Note that fairness can be simultaneously ensured across multiple sensitive attributes (e.g., ethnicity *and* political party) – the number of attributes a single candidate can have is captured by Δ . For example, in the New Zealand parliamentary election [Zealand, 1986] the parliament is required to include sufficient representation across 3 types of attributes: political parties, special interest groups, and Maori representation. Each candidate has an identity under each attribute, i.e., $\Delta = 3$.

This type of constraints generalize notions of proportionality that have arisen in the voting literature such as fully proportional representation [Monroe, 1995], by letting $\ell_i = \lfloor k \cdot n_i / n \rfloor$ and $u_i = \lceil k \cdot n_i / n \rceil$ for all i . (Here, n_i is the number of voters who prefer type i .) Similarly, one can ensure other notions such as degressive proportionality [Koriyama *et al.*, 2013] (e.g., satisfying Penrose’s square root law [Penrose, 1946]) and flexible proportionality [Brill *et al.*,]. In particular, the percent of representation need not be exactly specified, rather one can input an allowable range.

This is also general enough to ensure the outcome satisfies existing notions of fairness, such as disparate impact, statistical parity, and risk difference. For example, consider the case of groups that form a partition, and let m_i denote the number of voters that have type i . Given some $\xi_i \in [0, 1]$, for each P_i , we say committee of size k satisfies ξ -statistical parity if $||S \cap P_i|/k - m_i/n| \leq \xi_i$ (see [Dwork *et al.*, 2012] for the original definition). We can set fairness constraints that guarantee ξ -statistical parity by setting ℓ_i and u_i such that $\xi_i \geq 1 - \max\{| \ell_i/k - m_i/n |, | u_i/k - m_i/n |\}$ for all i .

	$\Delta \geq 3$	$\Delta = 1$	only u_i	only ℓ_i	$p = O(1)$	Unconstrained
A	$(1 - 1/e - o(1))$ -bi (Thm. 11)	$1 - 1/e$ (Thm. 14)	$\frac{1}{\Delta+1}$ -bi (Thm. 12)	$1 - 1/e - o(1)$ (Thm. 13)	$1 - 1/e$ (Full)	$1 - 1/e$ (NWF78a)
C	Feasibility NP-hard (Thm. 6-8)	$1 - 1/e + \varepsilon$	$O(\log \Delta / \Delta)$ (Thm. 9)	Feasibility NP-hard (Thm. 7)	$1 - 1/e + \varepsilon$	$1 - 1/e + \varepsilon$ (NW78)

Table 1: A summary of the results for the constrained monotone submodular (MS) multiwinner voting problem; each column denotes a different kind of group structure of the attributes. In Row “A” (for algorithm), the entry “ θ ” means that there exists a θ -approximation algorithm; “ θ -bi” means that the algorithm produces a θ -approximate solution compared to the optimal solution but the fairness constraints may be violated by a small multiplicative factor. In Row “C” (for complexity), the entry “ θ ” means that under the assumption $P \neq NP$, there does not exist a polynomial time algorithm that will always find a solution whose value is at most a θ factor from the optimal solution; “Feasibility NP-hard” means it is NP-hard to check whether there is a feasible solution satisfying all fairness constraints. $\varepsilon > 0$ is an arbitrary constant. Each entry is either a reference for the result, a theorem (Thm. i) in this paper, or will appear in the full version (Full) of this paper.

Voting Rules	$\Delta \geq 3$	$\Delta = 2$	$\Delta = 1$	$p = O(1)$	Unconstrained
SNTV	A $(1 - o(1))$ -bi (Thm. 11) C Feasibility NP-hard (Thm. 6-8)	P (Full) P	P (Full) P	P (Full) P	P P
α -CC	A $(1 - 1/e - o(1))$ -bi (Thm. 11) C Feasibility NP-hard (Thm. 6-8)	$(1 - 1/e - o(1))$ -bi (Thm. 11) $1 - 1/e + \varepsilon$ (SFS15)	$1 - 1/e$ (Thm. 14) $1 - 1/e + \varepsilon$ (SFS15)	$1 - 1/e$ (Full) $1 - 1/e + \varepsilon$ (SFS15)	$1 - 1/e$ (LB11) $1 - 1/e + \varepsilon$ (SFS15)
β -CC	A $(1 - 1/e - o(1))$ -bi (Thm. 11) C Feasibility NP-hard (Thm. 6-8)	$(1 - 1/e - o(1))$ -bi (Thm. 11) $1 - 1/e + \varepsilon$ (Full)	$1 - 1/e$ (Thm. 14) $1 - 1/e + \varepsilon$ (Full)	$1 - 1/e$ (Full) $1 - 1/e + \varepsilon$ (Full)	PTAS (SFS15) NP-hard (PRZ08)

Table 2: A summary of our results for the constrained monotone submodular (MS) multiwinner voting problem using three variants of the Chamberlin-Courant rule; each column denotes a different kind of group structure of the attributes. The definitions of (Thm. i), (Full), “A”, “C”, “ θ -bi”, “Feasibility NP-hard” and ε are the same as in Table 1. “P” means there exists a polynomial time exact algorithm.

Finally, we note that the fair multiwinner voting rule that results after adding fairness constraints continues to satisfy many nice properties (e.g., consistency, monotonicity, and fair variants of weak unanimity or committee monotonicity; see [Elkind *et al.*, 2014] for formal definitions) of the (unconstrained) voting rule. We refer the reader to the full version of this paper for the details [Celis *et al.*, 2017b].

3 Preliminaries

Now we present the formal definition of our model and the definitions of three monotone submodular voting rules: SNTV, α -CC and β -CC, which we consider as special cases.

Definition 1 (Our model: constrained multiwinner voting) *We are given a set C of m candidates, n voters together with $R = \{\succ_i\}_{i \in [n]}$ where \succ_i is the preference list (potentially incomplete and/or non-strict) over the m candidates for voter i , and a score function $\text{score}_R : 2^{[m]} \rightarrow R_{\geq 0}$ with an evaluation oracle, a desired number $k \in [m]$ of winners, arbitrary groups $P_1, \dots, P_p \subseteq C$ and integers $\ell_1, \dots, \ell_p, u_1, \dots, u_p \in \mathbb{Z}_{\geq 0}$. Given a size- k committee S , define the fairness constraints by $\ell_i \leq |S \cap P_i| \leq u_i, \forall i \in [p]$.*

Let $\mathcal{B} \subseteq 2^{[m]}$ denote the family of all size- k committees that satisfy all fairness constraints. The goal of the constrained multiwinner voting problem is to select an $S \in \mathcal{B}$ that maximizes $\text{score}_R(S)$. If the score function is monotone submodular,¹ we call the problem the constrained MS multiwinner voting problem.

If R is clear from the context, we denote score_R by score . This succinct description of \mathcal{B} and score allows us to design fast algorithms despite the fact that \mathcal{B} can be exponentially.

For a preference order \succ and a candidate $c \in C$, we write $\text{pos}_{\succ}(c)$ to denote the position of c in \succ (candidate ranked first has position 1 and ranked last has position m). Given a size- k committee $S \subseteq C$, we denote $\text{pos}_{\succ}(S) := (i_1, \dots, i_k)$ to be the sequence of positions of the candidates in S sorted

¹Recall that a function $f : 2^{[m]} \rightarrow R_{\geq 0}$ is a monotone submodular (MS) function if $f(A \cup B) + f(A \cap B) \leq f(A) + f(B)$ for all $A, B \subseteq [m]$ and $f(A) \leq f(B)$ for all $A \subseteq B$.

in increasing order with respect to \succ . Define $[m]_k$ to be the set of all size- k increasing sequences of elements from $[m]$.

Definition 2 (CC) *In the Chamberlin-Courant rule, there exists a positional score function $\gamma_m : [m] \rightarrow \mathbb{R}$ satisfying that $\gamma_m(i) \geq \gamma_m(j)$ if $1 \leq i < j \leq m$. Define $\gamma_{m,k}(i_1, \dots, i_k) = \max_{j \in [k]} \gamma_m(i_j) = \gamma_m(i_1)$ for any $(i_1, \dots, i_k) \in [m]_k$. The total score function is defined by $\text{score}(S) := \sum_{i \in [n]} \gamma_{m,k}(\text{pos}_{\succ_i}(S))$.*

Now we define three CC rules: SNTV, α -CC and β -CC.

Definition 3 *The SNTV rule uses the following positional score function: $\gamma_m(1) = 1$ and $\gamma_m(i) = 0$ for $i > 1$. Observe that SNTV only requires that each \succ_i includes the most preferred candidate of voter i .*

Definition 4 *The α -CC rule uses the following positional score function: $\gamma_m(i) = 1$ for $i \leq k$ and $\gamma_m(i) = 0$ for $i > k$. Observe that α -CC only requires that each \succ_i includes the top k preferred candidates (without ordering) of voter i .*

Definition 5 *The β -CC rule uses the following positional score function: $\gamma_m(i) = m - i$.*

4 Related Work

The study of total score functions and their resulting optimization problems have received much attention in recent years. Often the optimization problem turns out to be NP-hard; both α -CC and β -CC are NP-hard [Procaccia *et al.*, 2008], $1 - 1/e$ is the best approximation ratio for α -CC [Skowron *et al.*, 2015a], the Monroe rule is computationally hard even if the voting parameters are small [Betzler *et al.*, 2013], and the OWA-based rules are hard in general [Skowron *et al.*, 2015b] as are the goalbase rules in various settings [Uckelman *et al.*, 2009]. Hence, one must largely resort to developing approximation algorithms for these problems. Towards this, there has been a rich line of work [Lu and Boutilier, 2011; Skowron *et al.*, 2015a; 2015a; 2015b]. The majority of score functions for multiwinner voting rules that have been studied are monotone submodular. Algorithm design for such score functions have benefited from theoretical developments in the area of monotone

submodular function maximization [Nemhauser *et al.*, 1978; Calinescu *et al.*, 2011].

Some recent work also considers aspects of fairness in voting. Goalbase score functions, which specify an arbitrary set of logic constraints and let the score capture the number constraints satisfied [Uckelman, 2010; Uckelman *et al.*, 2009], could be used to ensure fairness. However, there are no known efficient algorithms to solve goalbase functions. Some recent literature studies single-winner voting in the multi-attribute setting; see the survey of Lang and Xia [2016]. The bi-apportionment model can handle up to two attributes (often political party and district) [Serafini and Simeone, 2012; Lari *et al.*, 2014]. Another related model is called constraint approval voting (CAP), with constraints on the numbers of winners from different categories of candidates, proposed by Brams [1990] and Potthoff [1990]. However, there is no efficient algorithm since the input in CAP is exponentially large in the number of attributes. Lang and Skowron [2016] also consider the problem of committee selection with multiple partitions; however, the goal is to produce a committee close to a given target composition of attributes as opposed to maximizing a score function.

More generally, our results contribute to the growing set of algorithms that incorporate fairness constraints to counter algorithmic bias in fundamental algorithmic problems such as classification [Dwork *et al.*, 2012; Zemel *et al.*, 2013; Zafar *et al.*, 2017b; 2017a], sampling [Celis *et al.*, 2016; 2017a; 2018a], ranking [Celis *et al.*, 2018b; Yang and Sotyanovich, 2017] and personalization [Celis and Vishnoi, 2017]. Independently of this paper, [Brederek *et al.*, 2018] also propose a model for multiwinner voting with a type of fairness constraints (referred to as *diversity constraints*).

5 Hardness Results

Here we outline hardness results for the constrained multiwinner voting problem; we refer the reader to the full version of this paper for the formal proofs [Celis *et al.*, 2017b]. We first address the complexity of the feasibility problem. Recall that Δ is the maximum number of groups in which a candidate can be. When a candidate may be part of 3 or more groups, just the feasibility problem can become NP-hard; even under mild feasibility conditions.

Theorem 6 (*NP-hardness of feasibility: $\Delta \geq 3$*) *The constrained multiwinner voting feasibility problem is NP-hard for any $\Delta \geq 3$, even if all $\ell_i = 0$ or all $u_i = |P_i|$.*

The proof uses a reduction from two NP-hard problems Δ -hypergraph matching problem [Hazan *et al.*, 2003] and the 3-regular vertex cover problem [Alimonti and Kann, 1997].

Furthermore, the next two theorems show that the feasibility problem remains hard even if one allows the size of the committee or the fairness constraints to be violated.

Theorem 7 (*Hardness of feasibility with committee size violation*) *Let $u_i = |P_i|$ for all $i \in [p]$. For any $\varepsilon > 0$, the following gap version of the constrained multiwinner voting feasibility problem is NP-hard: 1) Output YES if the input instance is feasible. 2) Output NO if there is no feasible solution of size less than $(1 - \varepsilon)k \ln p$.*

The proof uses a reduction from the constrained set multi-cover problem, which is hard to approximate [Berman *et al.*, 2004].

Theorem 8 (*Hardness of feasibility with fairness violations*) *Assume $\ell_i = 0$ for each $i \in [p]$. For every $\theta > 1$, the following violation variant of the constrained multiwinner voting feasibility problem is NP-hard: 1) Output YES if the input instance is feasible. 2) Output NO if there is no solution S of size k such that $|S \cap P_i| \leq \theta u_i$ for all $i \in [p]$.*

The proof relies on an inapproximability result for independent set [Zuckerman, 2006].

Lastly, we show that even if the constrained multiwinner voting instance is feasible, the hardness does not go away.

Theorem 9 (*Inapproximability for feasible instances*) *A feasible constrained multiwinner voting problem that satisfies $\ell_i = 0$ for each $i \in [p]$ is NP-hard to approximate within a factor of $\Omega(\log \Delta / \Delta)$.*

The reduction is from maximum Δ -hypergraph matching, which is hard to approximate within a factor of $\Omega(\log \Delta / \Delta)$ [Hazan *et al.*, 2003] and works even when the scoring function is the SNTV rule.

6 Algorithmic Results

Here we present our algorithmic results for the constrained multiwinner voting problem and sketch the proofs; we refer the reader to the full version of this paper for the complete proofs [Celis *et al.*, 2017b]. We first recall the notion of “multilinear extension”.

Definition 10 (*Multilinear extension*) *Given a monotone submodular function $f : \{0, 1\}^m \rightarrow \mathbb{R}_{\geq 0}$, the multilinear extension $F : [0, 1]^m \rightarrow \mathbb{R}_{\geq 0}$ is defined as follows: For $y = (y_1, \dots, y_m) \in [0, 1]^m$, denote \hat{y} to be a random vector in $\{0, 1\}^m$ where the j th coordinate is independently rounded to 1 with probability y_j or 0 otherwise. Then we let $F(y) = \mathbb{E}[f(\hat{y})] = \sum_{R \subseteq [m]} f(R) \prod_{i \in R} y_i \prod_{j \notin R} (1 - y_j)$.*

Define $\mathcal{B} \subseteq 2^{[m]}$ as the family of all committees of size k that satisfy the fairness constraints. Denote the polytope $B := \{y \in [0, 1]^m \mid \sum_{i \in [m]} y_i = k; \ell_j \leq \sum_{i \in P_j} y_i \leq u_j, \forall j \in [m]\}$ to be the set of all vectors that satisfy the cardinality constraint and all fairness constraints. Let $F : [0, 1]^m \rightarrow \mathbb{R}_{\geq 0}$ denote the multilinear extension of the total score function score. Let OPT be the optimal score of the constrained MS multiwinner voting problem.

6.1 The Case of $\Delta \geq 2$

Theorem 8 implies that it may be hard to find a committee only violating the fairness constraints by a small amount when $\Delta \geq 3$. The following theorem shows that a constant approximation solution can be achieved that violates all fairness constraints by at most a multiplicative factor for feasible instances. By taking the violation factors into account, the desired fairness can be achieved by setting tighter constraints.

Theorem 11 (*Bi-criterion algorithm when $\Delta \geq 2$*) *Consider a feasible constrained MS multiwinner voting instance with $OPT \gg \Omega(1)$. Let $L := \min_{i \in [p]} \ell_i$ and $U := \max_{i \in [p]} u_i$. Assume $2\sqrt{\ln p} / \sqrt{U} \leq 1$. There exists a randomized polynomial-time algorithm that outputs a committee S of size k with score at least $(1 - 1/e - o(1))OPT$ with constant probability, and S satisfies the following for all $i \in [p]$:*

$$(1 - 2\sqrt{\ln p} / \sqrt{L}) \ell_i \leq |S \cap P_i| \leq (1 + 2\sqrt{\ln p} / \sqrt{U}) u_i. \quad (1)$$

The approximation ratio is $1 - o(1)$ for the SNTV rule.

Before presenting the proof sketch, we discuss the assumption and the consequences of Theorem 11. First, the assumption $OPT \gg \Omega(1)$ is reasonable for several voting rules, such as the CC rule, the OWA-based rule, and the Monore rule. This is due to the fact that, if *enough* (e.g., at least $n/10$) voters have at least one representative in the optimal committee, then the total score of these rules is at least $n/10$.

Under reasonable assumptions, the violation in fairness constraints in the above theorem can be seen to be small. First, assume that no group is too small: $|P_i| \geq 0.15m \quad \forall i \in [m]$. Groups corresponding to gender, ethnicity and political opinions are often large. Combining this with the following conditions: $k \gg 100 \ln p$, and $\ell_i \approx k(|P_i|/m - 0.05)$, $u_i \geq k(|P_i|/m + 0.05)$ for all $i \in [p]$, we observe that $2\sqrt{\ln p}/\sqrt{L} \ll 0.66$ is a small number. Similarly, we can check that $2\sqrt{\ln p}/\sqrt{U} \ll 0.45 = \Omega(1)$ is also a small number. Thus, the violation of the group-fairness condition by Theorem 11 is small. We expect such algorithmic solutions to be deployed for the development of automated systems, such as movie selection on the airplane and news recommendation for websites, for which the above assumptions are natural.

Proof Sketch [of Theorem 11] We first obtain a fractional solution $y \in B$ by the continuous greedy algorithm in [Calinescu *et al.*, 2011, Section 3.1]. It follows that $F(y) \geq (1 - 1/e)OPT \gg \Omega(1)$ by [Calinescu *et al.*, 2011, Appendix A]. Next, we run the randomized swap rounding algorithm in [Chekuri *et al.*, 2010] and obtain a size- k committee S . By [Chekuri *et al.*, 2010, Theorem 2.1], we have $\mathbb{E}[\text{score}(S)] \geq F(y) \geq (1 - 1/e)OPT$, and for any $i \in [p]$ and any $\delta_1, \delta_2 > 0$, $\Pr[|S \cap P_i| \leq (1 - \delta_1)\ell_i] \leq e^{-\ell_i \delta_1^2/2}$, $\Pr[|S \cap P_i| \geq (1 + \delta_2)u_i] \leq (e^{\delta_2/(1+\delta_2)^{1+\delta_2}})^{u_i}$.

Let $\delta_1 = 2\sqrt{\ln p}/\sqrt{L}$ and $\delta_2 = 2\sqrt{\ln p}/\sqrt{U}$. Then by applying the union bound, we can show that S satisfies Equation (1) with probability at least $1 - 2/\sqrt{p}$. On the other hand, we have for any $\delta > 0$, $\Pr[\text{score}(S) \leq (1 - \theta)F(y)] \leq e^{-F(y)\theta^2/8}$ by [Chekuri *et al.*, 2010, Theorem 2.2]. Let $\theta = 0.01$. Since $F(y) \gg \Omega(1)$, we have $\text{score}(S) \geq (1 - 1/e - 0.01)OPT$ with probability $1 - o(1)$. \square

We now present the algorithmic results when there are only upper/lower fairness constraints. We first consider the case that all $\ell_i = 0$, i.e., each group is only required not to be over-represented. We call an algorithm (γ, θ) -approximation ($\gamma \in [0, 1], \theta \geq 1$) if the algorithm outputs a γ -approximate solution S such that $|S \cap P_i| \leq \theta \cdot u_i, \forall i \in [p]$.

Theorem 12 (Bi-criterion algorithm when $\ell_i = 0$) *Consider the constrained MS multiwinner voting problem satisfying $\ell_i = 0$ for all $i \in [p]$. Suppose we have a feasible solution \hat{S} in advance. For any $0 < \varepsilon < 1$, the following claims hold: 1) There exists a $(1/\Delta+1, 2)$ -approximation algorithm that runs in $O(mk/\Delta)$ time. 2) Suppose $u_i \geq 6 \log p/\varepsilon^2$ and $k \geq 6 \log p/\varepsilon^2$. There exists a polynomial-time $(1 - 1/e - O(\varepsilon), 2)$ -approximation algorithm.*

The assumption that we have a feasible solution \hat{S} is often satisfied and does not add computational overhead under natural

conditions on the data. For instance, let a_i denote the number of candidates that belong exclusively to group i . Then, if $a_i > \ell_i$ for all i and $\sum_i \min\{a_i, u_i\} > k$, a feasible solution exists and can be found in linear time via a greedy algorithm.

Proof Sketch [of Theorem 12] Define B' to be the collection of committees that have at most k candidates and satisfy all fairness constraints. For the first claim, we can prove that (C, B') is a Δ -extendible system since score is monotone submodular, we reduce the problem of finding $\arg \max_{S \in B'} \text{score}(S)$ to the monotone submodular maximization problem with Δ -extendible system, which has a $(1/\Delta+1)$ -approximation algorithm in $O(mk/\Delta)$ time by [Feldman *et al.*, 2017, Theorem 1]. Therefore, we can compute a committee $S_1 \in B'$ with $\text{score}(S_1) \geq (1/\Delta+1)OPT$ in $O(mk/\Delta)$ time by [Feldman *et al.*, 2017, Theorem 1]. Then from $\hat{S} \setminus S_1$, we arbitrarily select a set $S_2 \subseteq \hat{S}$ of $k - |S_1|$ candidates. We can verify that $S_1 \cup S_2$ is a $(1/\Delta+1, 2)$ -approximation solution.

For the second claim, we compute $y \in (1 - \varepsilon)B' := \{S \subseteq C : |S| \leq (1 - \varepsilon)k; |S \cap P_i| \leq (1 - \varepsilon)u_i, \forall i \in [p]\}$. Then we round y to a committee S_1 using the swap randomized rounding procedure in [Chekuri *et al.*, 2010]. By [Chekuri *et al.*, 2010, Theorem 5.2], S_1 can be guaranteed to be in B' , and $\text{score}(S_1) \geq (1 - 1/e - O(\varepsilon))OPT$. Then by the same argument as in the first claim, we can obtain a $(1 - 1/e - O(\varepsilon), 2)$ -approximation solution. \square

We now present the result for the case when we only have lower bound constraints, i.e., the upper bound constraints $u_i = |P_i|$ for all $i \in [p]$. Further, we assume that a committee of size $o(k/\ln \Delta)$ satisfying all fairness constraints exists, which is reasonable since, in practice, Δ roughly represents the number of attributes, like gender and ethnicity. Though there may exist many groups, the number of attributes is usually limited.

Theorem 13 (Approximation algorithm when $u_i = |P_i|$) *Given a constrained multiwinner voting instance satisfying $u_i = |P_i|$ for all $i \in [p]$ with a promise that a size- $o(k/\ln \Delta)$ committee satisfying all fairness constraints exists, there exists a $(1 - 1/e - o(1))$ -approximation algorithm.*

6.2 The Case of $\Delta = 1$

In this section, we consider the case $\Delta = 1$, i.e., $P_i \cap P_j = \emptyset$ for all $1 \leq i < j \leq p$ and prove the following theorem.

Theorem 14 (Algorithm for $\Delta = 1$) *Given a feasible constrained MS multiwinner voting instance with $\Delta = 1$, there exists a randomized polynomial-time $(1 - 1/e)$ -approximation algorithm.*

Algorithm DegreeOne.

1. Compute a fractional solution $y \in B$ by the continuous greedy algorithm in [Calinescu *et al.*, 2011, Section 3.1].
2. For $j \in [p]$, iteratively do the following until P_j has at most 1 non-integral coordinate: Arbitrarily select two candidates $c_i, c_{i'} \in P_j$ with fractional coordinates $y_i, y_{i'}$. Let $\delta_1 = \min\{1 - y_i, y_{i'}\}$ and $\delta_2 = \min\{y_i, 1 - y_{i'}\}$. Let $e_i = (0, \dots, 0, 1, 0, \dots, 0) \in \{0, 1\}^m$. Construct two vectors $y_1 = y + \delta_1(e_i - e_{i'})$ and $y_2 = y + \delta_2(e_{i'} - e_i)$. Let $y \leftarrow y_1$ with probability $\delta_2/\delta_1 + \delta_2$ and $y \leftarrow y_2$ otherwise.




Figure 1: Depicts the winners for the unconstrained (left) and prop w.r.t. voters (right) Bloc rule from all 1000 experiments; the constraints force all four quadrants to get the same number of winners.

3. W.l.o.g., assume y_1, \dots, y_γ are the remaining fractional coordinates. Iteratively do the same procedure as in Step 2 until y becomes an integral solution.
4. Output S whose indicator vector is $y = \mathbf{1}_S$.

Proof Sketch [of Theorem 14] It is not hard to verify the feasibility of S and the running time. For the approximation ratio, we have $F(y) \geq (1 - 1/e)OPT$ in Step 1 by [Calinescu *et al.*, 2011, Appendix A]. Then by [Calinescu *et al.*, 2011, Section 2.1], we can prove that $\mathbb{E}[F(y)] \geq F(y)$ for each iteration of Step 2 and 3. Thus, we conclude that $\mathbb{E}[\text{score}(S)] \geq F(y) \geq (1 - 1/e)OPT$. \square

7 Empirical Results

We compare the performance of winning committees without fairness constraints and with fairness constraints, under several commonly used MS multiwinner voting rules. We assume that the voter preferences are generated according to a two-dimensional Euclidean model as in [Schofield, 2007] which suggests that voters’ political opinions can be described sufficiently well in two dimensions. Multiwinner voting rules may introduce or exacerbate bias in such preference models. We show that the fairness constraints can prevent this, and moreover, the score attained by the fair result is close to the score attained by the optimal unconstrained (and hence, biased) committee.

7.1 Setup

Voting Rules. We consider five MS multiwinner voting rules: SNTV, Bloc, k -Borda, α -CC and β -CC. The definitions of SNTV, α -CC and β -CC appear in Section 3. Bloc outputs the k candidates that have the most voters listing them in the top k of their list. k -Borda outputs the k candidates with the largest sum of the Borda scores (Definition 5) that she receives from all voters. As a baseline, we also consider the Random voting rule that simply selects a committee uniformly at random. In all cases, we let $k = 12$ be the size of the desired committee.

Sampling Candidates, Voters and Preferences. We generate 400 voters and 120 candidates where each voter and candidate is represented by a point in the $[-3, 3] \times [-3, 3]$ square. $1/4$ of the voters are sampled uniformly at random from each quadrant, and $1/3$ of the candidates are sampled uniformly from the first quadrant, $1/4$ the second quadrant, $1/6$ the third quadrant, and $1/4$ the fourth quadrant. As in [Elkind *et al.*, 2017], we use Euclidean preferences; given a pair of

candidates $c_i, c_j \in \mathbb{R}^2$, a voter $a \in \mathbb{R}^2$ prefers c_i to c_j if $d(c_i, a) < d(c_j, a)$ where $d(\cdot, \cdot)$ is the Euclidean distance.

Groups and Fairness Constraints. We consider each quadrant to be a different group, and let P_i be the collection of candidates in the i -th quadrant. We compare the performance of three different types of constraints to the unconstrained optimal solution and the random baseline.

- **Proportional w.r.t. voters:** the number of winners is proportional to the number of voters in each quadrant.
- **Proportional w.r.t. candidates:** The number of winners is proportional to the number of candidates in each quadrant.
- **Relax:** The number of winners in each quadrant is allowed to be anywhere in the range between what would be proportional to voters and proportional to candidates.

Metrics. As a metric of fairness, we consider the Gini index, a well-studied metric for inequality: Let p be the total number of groups, and let n_i be the number of winners in each group P_i , then the Gini index is defined to be $\frac{\sum_{i=1}^p \sum_{j=1}^p |n_i - n_j|}{2p \sum_{i=1}^p n_j}$. This measures, on a scale from 0 to 1, how disproportionate the distribution of winners is amongst the groups (here, quadrants), with 1 meaning complete inequality and 0 meaning complete equality. In this context, the Gini index should be interpreted relationally, observing how different methods and constraints change the index up or down, as opposed to thinking of it as prescribing a correct outcome.

In addition, for each voting method we measure the price of fairness, i.e., the ratio between its score and the optimal unconstrained score. A ratio close to 100% indicates that despite adding fairness constraints the method attains close to an optimal score.

7.2 Results

We report the mean and the standard deviation of the Gini index, and the mean of the score ratio of 1000 repetitions in Table 3, and, as an example, depict the outcome of 1000 repetitions for the Bloc rule in Figure 1.

Unconstrained rules introduce bias. We observe that there can be bias in the unconstrained voting rules as the Gini index is large; in particular, for several rules, the inequality is larger than that of a random set of candidates suggesting that underlying disproportionalities in the set of candidates can be exacerbated by certain rules. The intuition is that having many candidates in the same quadrant thins their supporters due to competition, leading to disproportionately fewer representatives as winners.

For α -CC and β -CC, the unconstrained Gini index is close to that the Gini index when constraints are placed to be proportional with respect to candidates. However, the standard deviation in both cases is 0.06, which suggests that there can be significant bias in some outcomes of this process. The constrained process, on the other hand, will always satisfy the corresponding Gini index exactly.

The fairness constraints allow us to fix a particular distribution of winners (and hence corresponding Gini index), in this case either proportional with respect to voters or candidates, and hence allows us to de-bias the result in any desired manner. The result of the relaxed constraints, which allow the

Rule	Unconstrained		Prop w.r.t. voters		Prop w.r.t. candidates		Relax		Random	
	Gini	% Opt	Gini	% Opt	Gini	% Opt	Gini	% Opt	Gini	% Opt
SNTV	0.24 (0.09)	100	0 (0)	97.0	0.125 (0)	94.2	0.01 (0.02)	97.0	0.22 (0.09)	37.1
Bloc	0.28 (0.10)	100	0 (0)	91.6	0.125 (0)	88.4	0.00 (0.00)	91.6	0.22 (0.09)	61.9
k -Borda	0.24 (0.09)	100	0 (0)	98.9	0.125 (0)	99.3	0.11 (0.04)	99.3	0.22 (0.09)	72.6
α -CC	0.15 (0.06)	100	0 (0)	100	0.125 (0)	100	0.10 (0.05)	100	0.22 (0.09)	73.5
β -CC	0.11 (0.06)	100	0 (0)	100	0.125 (0)	100	0.07 (0.06)	100	0.22 (0.09)	95.8

Table 3: Under different rules, the average Gini index, the standard deviation of Gini index, and the average ratio of the constrained optimal score over the unconstrained optimal score.

distribution of winners within these ranges vary with some rules tending towards proportionality with respect to voters and others with respect to candidates.

The price of fairness is small. Because the feasible space of committees becomes smaller in the constrained settings, the optimal constrained score may be less than its unconstrained counterpart. However, in Table 3 we observe that the fairness constraints do not decrease the score by much; in fact for the α -CC or β -CC rules the score does not decrease at all.² This is not just a matter of there being many good committees; indeed a random committee in these settings does not perform well in comparison. Hence, the “price of fairness” is small in this setting; understanding this quantity more generally remains an important direction for future work.

8 Conclusion

In this paper, we propose a general model for multiwinner voting that can handle arbitrary MS score functions, can simultaneously ensure fairness across multiple sensitive attributes and can satisfy many different existing notions of fairness. We also develop efficient approximation algorithms for our model in different settings. The empirical results show that our model indeed eliminates bias compared to unconstrained voting, with a controlled price of fairness.

References

[Alimonti and Kann, 1997] Paola Alimonti and Viggo Kann. Hardness of approximating problems on cubic graphs. In *CIAC*, 1997.

[Aziz *et al.*, 2017a] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 2017.

[Aziz *et al.*, 2017b] Haris Aziz, Edith Elkind, Piotr Faliszewski, Martin Lackner, and Piotr Skowron. The Condorcet principle for multiwinner elections: From shortlisting to proportionality. In *IJCAI*, 2017.

[Berman *et al.*, 2004] Piotr Berman, Bhaskar DasGupta, and Eduardo D. Sontag. Randomized approximation algorithms for set multicover problems with applications to reverse engineering of protein and gene networks. In *RANDOM-APPROX*, 2004.

[Betzler *et al.*, 2013] Nadja Betzler, Arkadii Slinko, and Johannes Uhlmann. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research*, 2013.

²The ratio 100% for β -CC is due to rounding.

[Brams, 1990] Steven J Brams. Constrained approval voting: A voting system to elect a governing board. *Interfaces*, 1990.

[Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[Bredereck *et al.*, 2018] Robert Bredereck, Piotr Faliszewski, Ayumi Igarashi, Martin Lackner, and Piotr Skowron. Multiwinner elections with diversity constraints. In *AAAI*, 2018.

[Brill *et al.*,] Markus Brill, Jean-François Laslier, and Piotr Skowron. Multiwinner approval rules as apportionment methods. In *AAAI 2017*.

[Bundy, 2017] Alan Bundy. Preparing for the future of artificial intelligence. *AI Soc.*, 2017.

[Calinescu *et al.*, 2011] Gruiua Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 2011.

[Celis and Vishnoi, 2017] L. Elisa Celis and Nisheeth K. Vishnoi. Fair Personalization. *FATML*, 2017.

[Celis *et al.*, 2016] L. Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K Vishnoi. How to be fair and diverse? *FATML*, 2016.

[Celis *et al.*, 2017a] L. Elisa Celis, Amit Deshpande, Tarun Kathuria, Damian Straszak, and Nisheeth K. Vishnoi. On the complexity of constrained determinantal point processes. In *APPROX/RANDOM*, 2017.

[Celis *et al.*, 2017b] L. Elisa Celis, Lingxiao Huang, and Nisheeth K. Vishnoi. Multiwinner voting with fairness constraints. *CoRR*, abs/1710.10057, 2017.

[Celis *et al.*, 2018a] L. Elisa Celis, Vijay Keswani, Amit Deshpande, Tarun Kathuria, Damian Straszak, and Nisheeth K. Vishnoi. Fair and diverse DPP-based data summarization. In *ICML*, 2018.

[Celis *et al.*, 2018b] L Elisa Celis, Damian Straszak, and Nisheeth K Vishnoi. Ranking with fairness constraints. In *ICALP*, 2018.

[Chamberlin and Courant, 1983] John R Chamberlin and Paul N Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 1983.

[Chekuri *et al.*, 2010] Chandra Chekuri, Jan Vondrak, and Rico Zenklusen. Dependent randomized rounding via

- exchange properties of combinatorial structures. In *FOCS*, 2010.
- [Cox, 1994] Gary W Cox. Strategic voting equilibria under the single nontransferable vote. *American Political Science Review*, 1994.
- [Dwork *et al.*, 2012] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *ITCS*, 2012.
- [Elkind *et al.*, 2014] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. Properties of multiwinner voting rules. In *AAMAS*, 2014.
- [Elkind *et al.*, 2017] Edith Elkind, Piotr Faliszewski, Jean-François Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. What do multiwinner voting rules do? an experiment over the two-dimensional euclidean domain. In *AAAI*, 2017.
- [Faliszewski *et al.*, 2016] Piotr Faliszewski, Jean-François Laslier, Robert Schaefer, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Modeling representation of minorities under multiwinner voting rules. *CoRR*, abs/1604.02364, 2016.
- [Feldman *et al.*, 2017] Moran Feldman, Christopher Harshaw, and Amin Karbasi. Greed is good: Near-optimal submodular maximization via greedy optimization. In *COLT*, 2017.
- [Hazan *et al.*, 2003] Elad Hazan, Shmuel Safra, and Oded Schwartz. On the complexity of approximating k -dimensional matching. In *RANDOM-APPROX*. Springer, 2003.
- [Kay *et al.*, 2015] Matthew Kay, Cynthia Matuszek, and Sean A Munson. Unequal representation and gender stereotypes in image search results for occupations. In *CHI*, 2015.
- [Koriyama *et al.*, 2013] Yukio Koriyama, Antonin Macé, Rafael Treibich, and Jean-François Laslier. Optimal apportionment. *Journal of Political Economy*, 2013.
- [Lang and Skowron, 2016] Jérôme Lang and Piotr Krzysztof Skowron. Multi-attribute proportional representation. In *AAAI*, 2016.
- [Lari *et al.*, 2014] Isabella Lari, Federica Ricca, and Andrea Scozzari. Bidimensional allocation of seats via zero-one matrices with given line sums. *Annals OR*, 2014.
- [Lu and Boutilier, 2011] Tyler Lu and Craig Boutilier. Budgeted social choice: From consensus to personalized decision making. In *IJCAI*, 2011.
- [Monroe, 1995] Burt L Monroe. Fully proportional representation. *American Political Science Review*, 1995.
- [Nemhauser *et al.*, 1978] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions - I. *Mathematical Programming*, 1978.
- [Penrose, 1946] Lionel S Penrose. The elementary statistics of majority voting. *Journal of the Royal Statistical Society*, 1946.
- [Potthoff, 1990] R Potthoff. Use of linear programming for constrained approval voting. *Interfaces*, 1990.
- [Procaccia *et al.*, 2008] Ariel D Procaccia, Jeffrey S Rosenschein, and Aviv Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 2008.
- [Representation2020, 2017] Representation2020. http://www.representation2020.com/voting_rules_pose_barrier_for_women, 2017.
- [Schofield, 2007] Norman Schofield. *The spatial model of politics*. Routledge, 2007.
- [Serafini and Simeone, 2012] Paolo Serafini and Bruno Simeone. Parametric maximum flow methods for minimax approximation of target quotas in biproportional apportionment. *Networks*, 2012.
- [Skowron *et al.*, 2015a] Piotr Skowron, Piotr Faliszewski, and Arkadii Slinko. Achieving fully proportional representation: Approximability results. *Artificial Intelligence*, 2015.
- [Skowron *et al.*, 2015b] Piotr Krzysztof Skowron, Piotr Faliszewski, and Jérôme Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *AAAI*, 2015.
- [Uckelman *et al.*, 2009] Joel Uckelman, Yann Chevaleyre, Ulle Endriss, and Jérôme Lang. Representing utility functions via weighted goals. *Mathematical Logic Quarterly*, 2009.
- [Uckelman, 2010] Joel Uckelman. Alice and Bob will fight: the problem of electing a committee in the presence of candidate interdependence. In *ECAI*. IOS Press, 2010.
- [Yang and Stoyanovich, 2017] Ke Yang and Julia Stoyanovich. Measuring fairness in ranked outputs. In *SSDBM*, 2017.
- [Zafar *et al.*, 2017a] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *WWW*, 2017.
- [Zafar *et al.*, 2017b] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness constraints: Mechanisms for fair classification. In *AISTATS*, 2017.
- [Zealand, 1986] New Zealand. *Report of the Royal Commission on the Electoral System: Towards a Better Democracy*. Government Printer, South Africa, 1986.
- [Zemel *et al.*, 2013] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *ICML*, 2013.
- [Zuckerman, 2006] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *STOC*, 2006.