Multiwinner Rules on Paths From k-Borda to Chamberlin–Courant*

Piotr Faliszewski[1] and **Piotr Skowron**[2] and **Arkadii Slinko**[3] and **Nimrod Talmon**[4]
[1] AGH University, Krakow, Poland
[2] Institut für Softwaretechnik und Theoretische Informatik, TU Berlin, Berlin, Germany
[3] University of Auckland, Auckland, New Zealand
[4] Weizmann Institute of Science, Rehovot, Israel
faliszew@agh.edu.pl, p.k.skowron@gmail.com, a.slinko@auckland.ac.nz, nimrodtalmon77@gmail.com

Abstract

The classical multiwinner rules are designed for particular purposes. For example, variants of k-Borda are used to find k best competitors in judging contests while the Chamberlin-Courant rule is used to select a diverse set of k products. These rules represent two extremes of the multiwinner world. At times, however, one might need to find an appropriate trade-off between these two extremes. We explore continuous transitions from k-Borda to Chamberlin-Courant and study intermediate rules.

1 Introduction

The multiwinner voting rule k-Borda [Debord, 1992] and the rule of Chamberlin and Courant [1983] are based on very different ideas and they are used for very different purposes. For example, k-Borda is a maximum likelihood estimator [Procaccia *et al.*, 2012], which means that it works well when there is an objectively best committee and all the votes are noisy perceptions of the truth. Intuitively speaking, this means that we assume that votes "closer to the truth" are more likely than those "farther from the truth" and so we seek the area of the "highest density of the votes" to derive the winning committee from the votes there (ignoring outliers and voter groups of smaller density). One consequence of this approach is that the members in the winning committee under k-Borda tend to be very similar to each other (for a visual justification of these intuitive claims, we point the reader to the work of Elkind et al. [2017a] and to the histograms in our experimental section). Thus, k-Borda is very suitable for applications such as finding finalists of various competitions (based on the rankings provided by the judges)[1], or—generally speaking—for shortlisting tasks where the candidates are evaluated on a single criterion (Barberà and Coelho [2008] and Elkind et al. [2017b] offer discussions regarding shortlisting tasks).

The Chamberlin-Courant rule (the CC rule), on the other hand, was designed to determine as diverse set of alternatives as possible (e.g., a set of movies that an airline may offer on its long-range flights [Elkind *et al.*, 2017b]) so that every shade of preference existing in the society is catered for (Elkind et al. [2017a] justify this intuition visually; see also our experimental section). Consequently, the CC rule is very well suited for applications where we select a set of items from which each voter chooses one that he or she desires most (see the works of Lu and Boutilier [2011], Skowron et al. [2016a], and Faliszewski et al. [2016a] for further relevant scenarios in addition to choosing movies). However it is possible that some members of the Chamberlin-Courant elected committee will be liked by a minority of voters and will not score high on the excellence criterion.

To summarize, k-Borda and Chamberlin–Courant are very different rules and pursue different objectives: excellence versus diversity. These are two opposing objectives and the two rules are as far apart in the 'space of voting rules' as possible.

Nonetheless, there are many applications where neither excellence nor diversity are the most appropriate goals to aim for. For example, consider shortlisting candidates for an academic position in computer science. On the one hand, we may wish to shortlist simply the strongest candidates, but, on the other hand, we would also like to secure certain diversity among the selected candidates with respect to areas of specialization, gender balance, etc. Thus, while excellence is still a crucial criterion, we would also like to put significant weight on diversity (this is also related to proportional representation with respect to multiple external attributes [Lang and Skowron, 2016]).

Similarly, in the movie-selection example, while the desire to choose a diverse set of movies is our main driving force, usually we also want to make sure that more popular types of movies appear in greater numbers than the less popular ones (indeed, on a typical long-range flight we see more Hollywood blockbuster movies than, say, documentaries).

This means that many real-life, business-driven applications of multiwinner voting require rules that strike some sort of a compromise between excellence and diversity, sometimes putting more focus on the former and sometimes giving priority to the latter (see also the discussion of Ratliff and Saari [2014]). However, until recently, in the literature we have had a discrete set of voting rules and nothing in be-

*Piotr Faliszewski was supported by the National Science Centre, Poland, under project 2016/21/B/ST6/01509, Piotr Skowron was supported by a Humboldt Research Fellowship for Postdoctoral Researchers, Arkadii Slinko was supported by the Royal Society of NZ Marsden Fund UOA-254.

[1]See, e.g., the work of Gordon and Truchon [2008], where the problem of aggregating the judges' rankings in figure skating is treated from the maximum likelihood perspective.

tween them. In this paper we attempt to rectify this situation by introducing families of rules that, in a certain sense, lie on a "path between k-Borda and the Chamberlin–Courant rule". This is made possible due to the recently introduced class of committee scoring rules [Faliszewski *et al.*, 2016a] and by the fact that both k-Borda and the Chamberlin–Courant rule belong to this class. The committee scoring rules work as follows. To each possible committee, each voter assigns a score based on the positions of the committee members in his or her ranking. These scores are summed up and, as a result, each committee obtains a societal score. The committee with the highest societal score wins.

Both k-Borda and the Chamberlin-Courant rule use classical Borda scores to evaluate individual committee members. If the total number of alternatives is m and a candidate is ranked in the ith position in a voter's ranking, then this voter's Borda score of that candidate is $m - i$. For calculating the score of a committee, according to a given voter, under k-Borda we sum up all the scores of the committee members while Chamberlin-Courant takes the maximal score. For example, if the Borda scores of the committee members (from an individual voter) are $b_1 > b_2 > \ldots > b_k$, then the k-Borda score of this committee (from this voter) would be $b_1 + b_2 + \ldots + b_k$, while the Chamberlin-Courant score (from the same voter) would be b_1.

Skowron et al. [2016a] suggested to calculate the committee score as $\lambda_1 b_1 + \ldots + \lambda_k b_k$, where $\Lambda = (\lambda_1, \ldots, \lambda_k)$ is a vector of numbers between 0 and 1, referred to as the ordered weighted average operator (OWA); the rules that use OWA operators in calculating the scores of committees are called OWA-based. In particular, k-Borda and Chamberlin–Courant rules are OWA-based relative to the vectors $(1, \ldots, 1)$ and $(1, 0, \ldots, 0)$, respectively. We consider two families of OWA-based rules that connect these two extremes. The first is defined by vectors with t ones followed by $k - t$ zeros (where t ranges between 1 and k), and the second one uses vectors of the form $(1, 1/2^q, 1/3^q, \ldots, 1/k^q)$, for $q \in (0, \infty)$.

Another way to connect k-Borda and the Chamberlin-Courant rule is to observe that summing up the Borda scores of the committee members—as in k-Borda—simply means taking the ℓ_1-norm of the vector of those Borda scores, whereas taking the maximal score—as in CC—means taking the ℓ_∞-norm. Thus, we consider the family of rules between k-Borda and Chamberlin–Courant that use the ℓ_p-norms with $p \in [1, \infty)$ (we refer to these rules as ℓ_p-Borda rules).

We test our rules on elections from the 2D Euclidean domain, following Elkind et al. [2017a]. The rules behave in a rather unexpected way. For example, as we move from ℓ_1-Borda to ℓ_∞-Borda, we first observe a rapid phase transition from rules that are similar to k-Borda to rules that resemble the Bloc rule (a rule that, at first sight, does not seem to have anything to do with ℓ_p-Borda rules) and then there is a smooth transition to rules that are similar to CC. Rules using our first family of OWA vectors behave similarly (but "more smoothly", which is quite surprising given that these OWAs have a much more discrete nature), whereas the rules based on the second family of OWAs follow a very different path.

As a consequence of our studies, we also find more decisive variants of the k-Borda and Chamberlin–Courant rules.

2 Preliminaries

In this paper we focus on committee scoring rules defined by Elkind et al. [2017b]. In our model, an election $E = (C, V)$ consists of a set $C = \{c_1, \ldots, c_m\}$ of candidates and a collection $V = (v_1, \ldots, v_n)$ of voters, where each voter v_i has a preference order $\succ_i$, thereby ranking the candidates from the most desirable to the least desirable one. For a candidate c and a voter v, we write $\mathrm{pos}_v(c)$ to denote the position of c in v's preference order (e.g., $\mathrm{pos}_v(c) = 1$ if v ranks c in the top position and $\mathrm{pos}_v(c) = m$ if v ranks c last). We refer to subsets of C as committees and if S is a committee and v is a voter, then by $\mathrm{pos}_v(S)$ we mean the sequence obtained by sorting the set $\{\mathrm{pos}_v(c) \mid c \in S\}$ in the increasing order.

For positive integers m and k, we write $[m]$ to denote the set $\{1, \ldots, m\}$ and $[m]_k$ to denote the set of all length-k increasing sequences of elements from $[m]$. In other words, for a preference order over m candidates, $[m]$ is the set of all candidate positions in the voters' rankings and $[m]_k$ is the set of all possible positions that a committee of size k can occupy. For two committee positions $I = (i_1, \ldots, i_k)$ and $J = (j_1, \ldots, j_k)$ from $[m]_k$, we say that I *dominates* J, denoted $I \succeq J$, if for each $t \in [k]$, it holds that $i_t \leq j_t$.

A multiwinner rule $\mathcal{R}$ is a function that, given an election $E = (C, V)$ and a positive integer k, outputs a family $\mathcal{R}(E, k)$ of size-k committees that tie as winners of E (we typically expect a single winning committee, but there can be more due to symmetries in the preference orders).

A committee scoring function, for elections with m candidates and committees of size k, is a function $f_{m,k} \colon [m]_k \to \mathbb{R}$ that assigns a numerical score to each committee position, so that if $I, J \in [m]_k$ are two positions and $I \succeq J$, then $f(I) \geq f(J)$. For an election $E = (C, V)$ with m candidates and a size-k committee S, we define the $f_{m,k}$-score of S as:

$$\mathrm{score}_{f_{m,k}}(S, E) = \textstyle\sum_{v \in V} f_{m,k}(\mathrm{pos}_v(S)).$$

Definition 1. *Let $f = \{f_{m,k}\}_{1 \leq k \leq m}$ be a family of committee scoring functions. A committee scoring rule $\mathcal{R}_f$ is a multiwinner rule that, given an election $E = (C, V)$ and committee size k, outputs those committees S (of size k) for which $\mathrm{score}_{f_{|C|,k}}(S, E)$ is the highest.*

For m candidates, the (single-winner) Borda score awarded to a candidate for occupying position i in a ranking of some voter is $\beta_m(i) = m - i$. The k-Borda and Chamberlin–Courant committee scoring rules are defined through the following scoring functions:

1. For the k-Borda rule, the scoring functions are:
$$f^{k\text{-Borda}}_{m,k}(i_1, \ldots, i_k) = \beta_m(i_1) + \ldots + \beta_m(i_k).$$
To find a k-Borda winning committee it suffices to compute the Borda scores separately for each candidate and pick those k candidates with the highest scores. That is, k-Borda does not consider any dependencies between the candidates and simply picks the k best ones according to the Borda social welfare function which ranks all candidates by their Borda scores.

2. For the CC rule, the scoring functions are:
$$f^{\mathrm{CC}}_{m,k}(i_1, \ldots, i_k) = \beta_m(i_1).$$

Intuitively, under the CC rule each voter treats his or her highest-ranked committee member as his or her representative and the score of a committee is the sum of the Borda scores that the voters give to their representatives.

In our discussion, we will also refer to the Bloc rule:

3. Under Bloc, each voter names his or her k favorite candidates and the winning committee consists of those mentioned most frequently. In other words, Bloc uses the scoring functions
$$f^{\mathrm{Bloc}}_{m,k}(i_1,\ldots,i_k) = \textstyle\sum_{t=1}^{k} \alpha_k(i_t),$$
where $\alpha_k(i) = 1$ for $i \leq k$ and $\alpha_k(i) = 0$ otherwise.

Example 1. *Consider an election with the following votes:*
$$v_1\colon a \succ b \succ c \succ d \succ e \succ f \succ g \succ h,$$
$$v_2\colon f \succ h \succ e \succ g \succ d \succ c \succ b \succ a,$$
$$v_3\colon a \succ b \succ g \succ h \succ e \succ f \succ c \succ d,$$
$$v_4\colon a \succ c \succ d \succ g \succ h \succ b \succ f \succ e.$$
Under 4*-Borda, the unique winning committee of size* 4 *is* $\{a,b,c,g\}$*, under CC all committees that include* a *and* f *win.*

While winning committees for k-Borda and Bloc can be computed in polynomial time, the CC rule is computationally hard [Procaccia *et al.*, 2008; Lu and Boutilier, 2011; Betzler *et al.*, 2013]. Nonetheless, CC becomes polynomial-time solvable under certain domain restrictions [Betzler *et al.*, 2013; Skowron *et al.*, 2015b; Yu *et al.*, 2013; Peters and Elkind, 2016] and there are approximation algorithms and heuristics for it [Skowron *et al.*, 2015a; Faliszewski *et al.*, 2016c].

Naturally, many other committee scoring rules exist. We point the reader to the works of Faliszewski et al. [2016a; 2016b] for an overview of the internal structure of the class of committee scoring rules, and to the work of Skowron et al. [2016b] for their axiomatic characterization.

3 Three Paths Between *k*-Borda and CC

Below we describe the three families of rules that we view as paths between k-Borda and CC. We start with two families based on OWA operators, then we introduce the ℓ_p-Borda family of rules, and finally we discuss convergence of the rules in our families to CC and to k-Borda, as well as some of their properties.

The OWA-Based Paths. A family $\Lambda = (\Lambda^k)_{k\in\mathbb{N}}$ of OWA operators, $\Lambda^k = (\lambda^k_1,\ldots,\lambda^k_k)$, with one operator for each dimension, defines a committee scoring rule with the following scoring functions:
$$f^{\Lambda}_{m,k}(i_1,\ldots,i_k) = \lambda^k_1\beta_m(i_1) + \ldots + \lambda^k_k\beta_m(i_k). \qquad (1)$$
Multiwinner rules of this type were introduced by Skowron et al. [2016a] in a somewhat more general context (i.e., not focusing only on the Borda scores), and were later studied by Faliszewski et al. [2016a].[2]

We consider the following OWA operators Λ^k:

1. For each positive integer t, the t-Best family of OWA operators has vectors of t ones followed by $t-k$ zeros.
2. For each non-negative real number p, the p-Harmonic OWA operators are of the form $(1, {}^{1}\!/{}_{2^p}, {}^{1}\!/{}_{3^p}, \ldots, {}^{1}\!/{}_{k^p})$.

For each t, we set t-Borda to be the rule defined using Equation (1) and the t-Best OWA operators (note that this notation is consistent with the name k-Borda), and for each non-negative real number p, we let p-HarmonicBorda be defined analogously, but using the p-Harmonic OWAs. (1-Harmonic OWAs are also used in the definition of the PAV rule [Thiele, 1895; Kilgour and Marschall, 2012], that was shown to have very good axiomatic properties [Aziz *et al.*, 2017].)

Example 2. *Consider the election from Example 1. The* 2*-Borda rule, for* $k = 4$*, chooses committees* $\{a,b,f,h\}$ *and* $\{a,c,f,h\}$ *(both of which are winning under CC),* 2*-HarmonicBorda chooses only* $\{a,b,f,h\}$*, while* ${}^{1}\!/{}_{2}$*-HarmonicBorda chooses* $\{a,b,c,h\}$*, which differs from the* 4*-Borda winning committee only by replacing* g *with* h*.*

Intuitively, as p gets larger, p-HarmonicBorda becomes more similar to CC (and ${}^{1}\!/{}_{p}$-HarmonicBorda becomes more similar to k-Borda). This intuition is almost correct and later we will make it formal.

The rules on our paths (with the exception of k-Borda) are NP-hard to compute (this follows immediately from the results of Skowron et al. [2016a]).

Proposition 1. *For each positive integer* $t > 1$*, it is* NP*-hard to decide if a winning committee under* t*-Borda has at least a given score. The same holds for each positive real number* p *and* p*-HarmonicBorda.*

While the above result speaks of values of t that are constant (and, in particular, do not depend on k), Skowron et al. also showed that, e.g., $(k-1)$-Borda is NP-hard to compute.

The $\boldsymbol{\ell_p}$-Path. Let p be a real number, $p \geq 1$. By ℓ_p-Borda we mean the committee scoring rule defined by the committee scoring functions:
$$f^{\ell_p\text{-}\beta}_{m,k}(i_1,\ldots,i_k) = \sqrt[p]{\beta_m(i_1)^p + \ldots + \beta_m(i_k)^p}.$$
That is, under ℓ_p-Borda, the score associated with a committee position $(i_1,\ldots,i_k)$ is the ℓ_p-norm of the vector $(\beta(i_1),\ldots,\beta(i_k))$. Naturally, ℓ_1-Borda is k-Borda and ℓ_∞-Borda is CC. We are interested, however, in ℓ_p-Borda rules where p is not ∞.

Example 3. *Consider the election from Example 1.* ℓ_2*-Borda chooses committee* $\{a,b,c,f\}$*, a CC winning committee that is close to the* k*-Borda winning committee (the fact that the committee is winning under CC is mostly a coincidence). Note that the rule in this example and all the rules in Examples 1 and 2 give different sets of committees as outcomes.*

All the ℓ_p-Borda rules for $p > 1$ are NP-hard to compute.

Theorem 2. *For each rational* $p > 1$*, deciding whether there is a winning committee with at least a given score is* NP*-hard.*

In principle, the idea behind the proof of this theorem is close to that behind the proof of the NP-hardness of CC [Lu and Boutilier, 2011]. However, it requires significant technical work to introduce sufficient (but polynomially-bounded)

[2] See also the work of Aziz et al. [2017]. Goldsmith et al. [2014] and Elkind and Ismaili [2015] provide other ways in which OWA operators can be used in voting rules.

paddings that, roughly speaking, ensures that ℓ_p-Borda rules behave similarly to CC.

Convergence of the Rules on the Paths. While the t-Borda family of rules contains k-Borda and CC by definition, the situation is more intricate for the ℓ_p-Borda and p-HarmonicBorda families. In particular, we are interested in convergence of p-HarmonicBorda and ℓ_p-Borda to both k-Borda and CC. (Note that 0-HarmonicBorda and ℓ_1-Borda are simply nicknames for k-Borda, but the rules that could be dubbed $\lim_{p\to 0^+}$ p-HarmonicBorda and $\lim_{p\to 1^+}$ ℓ_p-Borda offer some advantages regarding decisiveness; due to limited space, we mostly focus on convergence to CC, however).

We next show how well the CC winning committees approximate the ℓ_p-Borda and p-HarmonicBorda winning ones.

Definition 2. *Let $\mathcal{R}_f$ be a committee scoring rule defined through a family of committee scoring functions $f = (f_{m,k})_{1\le k\le m}$. Let $E = (C, V)$ be an election and W be a winning committee of E under $\mathcal{R}_f$. Let α be a real number such that $0 < \alpha < 1$. We say that a committee S is an* α-approximate winner *for the $\mathcal{R}_f$-election E, if* $\mathrm{score}_{f_{|C|,k}}(S, E) \ge \alpha \cdot \mathrm{score}_{f_{|C|,k}}(W, E)$.[3]

Proposition 3. *Let E be an election with m candidates and let k be the target committee size.*

1. *For each $p \ge 1$, each CC winning committee for E is a $\frac{1}{\sqrt[p]{k}}$-approximate ℓ_p-Borda winner for E.*
2. *For each $p \ge 0$, each CC winning committee for E is a $1/(\sum_{j=1}^{k} 1/j^p)$-approximate p-HarmonicBorda winner for E.*

Proof. We consider the case of ℓ_p-Borda (the case of p-HarmonicBorda is analogous and omitted due to restricted space). Let W_p be an ℓ_p-Borda winning committee for E and let W_∞ be a CC winning committee. For brevity, we write f to refer to the scoring function $f^{\ell_p\text{-}\beta}_{m,k}$. For a given committee S and each voter v we have:

$$f(\mathrm{pos}_v(S)) \le \sqrt[p]{k} \cdot f^{CC}_{m,k}(\mathrm{pos}_v(S)) \le \sqrt[p]{k} \cdot f(pos_v(S))).$$

Thus it holds that:

$$\begin{aligned}
&\mathrm{score}_f(W_p, E) = \textstyle\sum_{v\in V} f(\mathrm{pos}_v(W_p))\\
&\le \sqrt[p]{k} \cdot \textstyle\sum_{v\in V} f^{CC}_{m,k}(\mathrm{pos}_v(W_p))\\
&\le \sqrt[p]{k} \cdot \textstyle\sum_{v\in V} f^{CC}_{m,k}(\mathrm{pos}_v(W_\infty))\\
&\le \sqrt[p]{k} \cdot \textstyle\sum_{v\in V} f(\mathrm{pos}_v(W_\infty)) = \sqrt[p]{k} \cdot \mathrm{score}_f(W_\infty, E),
\end{aligned}$$

which proves the statement. □

The estimates in Proposition 3 are asymptotically tight.

Example 4. *Let m, n, and k be three positive integers such that $m \ge 2k$. Consider an election E with candidates $c_1, \dots, c_m$ and with $n + (k-1)$ voters, where k is the target committee size. For each $i \in [k-1]$, the ith voter ranks c_{m-i+1} first; and each of the following n voters has preference order $c_1 \succ c_2 \succ \cdots \succ c_m$. For large enough n, we have that: (a) $c_1, \dots, c_k$ is the unique winning committee under each ℓ_p-Borda and each p-HarmonicBorda rule; (b) $c_1, c_m, c_{m-1}, \dots, c_{m-k+2}$ is the unique CC winning committee; (c) this committee is an $O(1/\sqrt[p]{k})$-approximate ℓ_p-Borda winner for E and an $O(1/(\sum_{j=1}^{k} \frac{1}{j^p}))$-approximate p-HarmonicBorda winner for E.*

[3]This definition relies on the particular scoring functions used.

For a given committee size, Proposition 3 indicates which values of p may lead to rules very similar to the CC rule.

Example 5. *Let us consider some election E and the committee size $k = 20$. CC winning committees are 0.54-approximations of ℓ_5-Borda winning committees (so, intuitively, we should not expect ℓ_5-Borda and CC results to be similar), while they are 0.94-approximations of ℓ_{50}-Borda winning committees (so we expect similarity). Similarly, CC winning committees are 0.62-approximations of 2-HarmonicBorda winning committees, but 0.96-approximation of 5-HarmonicBorda ones.*

Above, we studied convergence of the scores of the winning committees. Now we move to the convergence of actual winning committees (these issues differ as one could imagine two disjoint committees with nearly identical scores).

Definition 3. *Let $(\mathcal{R}_i)_{i\in\mathbb{N}}$ be a family of multiwinner voting rules. We say that this sequence* converges election-wise*, if for each election E and each target committee size k, there exists a positive integer i_0 such that $\mathcal{R}_{i_0}(E, k) = \mathcal{R}_i(E, k)$ for each $i \ge i_0$.*

A sequence $(\mathcal{R}_i)_{i\in\mathbb{N}}$ of rules that converges election-wise defines a new multiwinner rule: Given an election E and a target committee size k, it outputs the committee $\mathcal{R}_i(E, k)$ for sufficiently large i (when the sequence of outputs becomes stable).

A multiwinner rule $\mathcal{R}$ is a *refinement* of a rule $\mathcal{Q}$ if for each election E and each committee size k, $\mathcal{R}(E, k) \subseteq \mathcal{Q}(E, k)$.

Theorem 4. *The sequence of rules $(p\text{-HarmonicBorda})_{p\in\mathbb{N}}$ converges, as $p \to \infty$, to a refinement of CC. The sequence $(1/p\text{-HarmonicBorda})_{p\in\mathbb{N}}$ converges to a refinement of k-Borda.*

Proof. Let us consider the sequence $(p\text{-HarmonicBorda})_{p\in\mathbb{N}}$. We may assume that p changes continuously, i.e., that p is a real number. We fix an election E and the size of committee k. For a committee Z, we write $\mathrm{score}_p(Z)$ to denote its p-HarmonicBorda score in E. It suffices to show that for each two distinct committees X and Y, starting from some p_0 and for all $p > p_0$, we have either $\mathrm{score}_p(X) > \mathrm{score}_p(Y)$ or $\mathrm{score}_p(X) = \mathrm{score}_p(Y)$ or $\mathrm{score}_p(X) < \mathrm{score}_p(Y)$. Suppose $\mathrm{pos}_{v_i}(X) = \{\ell^{(i)}_1, \dots, \ell^{(i)}_k\}$. Then,

$$\begin{aligned}
\mathrm{score}_p(X) &= \textstyle\sum_{i=1}^{n}\sum_{j=1}^{k} \frac{1}{j^p}\beta_m(\ell^{(i)}_j)\\
&= \textstyle\sum_{j=1}^{k} \frac{1}{j^p}\sum_{i=1}^{n} \beta_m(\ell^{(i)}_j) = \sum_{j=1}^{k} \frac{1}{j^p} a_j.
\end{aligned}$$

where $a_j = \sum_{i=1}^{n} \beta_m(\ell^{(i)}_j)$. Similarly, $\mathrm{score}_p(Y) = \sum_{j=1}^{k} \frac{1}{j^p} b_j$ for some $b_1, \dots, b_k$. If $a_j = b_j$ for all $j \in [k]$, then the scores of X and Y are always the same, irrespective of the value of p. Suppose this does not happen.

As p grows, the relation between $\mathrm{score}_p(X)$ and $\mathrm{score}_p(Y)$ may change several times, but we claim that this number of changes is smaller than k. We will show that the difference

$$\mathrm{score}_p(X) - \mathrm{score}_p(Y) = \sum_{j=1}^{k} \frac{1}{j^p}(a_j - b_j)$$

cannot be equal to zero for k distinct values of p. Assume the contrary, i.e., that it is equal to zero for $p_1, \dots, p_k$. Then:

$$\begin{bmatrix} 1 & (1/2)^{p_1} & \cdots & (1/k)^{p_1} \\ 1 & (1/2)^{p_2} & \cdots & (1/k)^{p_2} \\ \cdots & \cdots & \cdots & \cdots \\ 1 & (1/2)^{p_k} & \cdots & (1/k)^{p_k} \end{bmatrix} \begin{bmatrix} a_1 - b_1 \\ a_2 - b_2 \\ \vdots \\ a_k - b_k \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

Since $a_j - b_j \neq 0$ for at least one $j \in [k]$, the determinant of the matrix on the left-hand side must be zero. However, this is a generalized Vandermonde matrix which is known to have a non-zero determinant [Wen and Zhang, 2006]. This completes the proof of our claim that the relations between the scores of committees are fixed for sufficiently large p.

As p approaches ∞, for any committee X we have $\mathrm{score}_p(X) \to \mathrm{score}_\infty(X)$; thus, the winning committees of the limiting voting rule have to be CC winning committees and it is a refinement of CC. The argument for k-Borda ($p \to 0$) is similar. □

The situation with ℓ_p-Borda rules is much more difficult and we only obtain a weaker result.

Proposition 5. *For each election E and each committee size k, there is a value p_0 such that for each $p > p_0$, ℓ_p-Borda(E,k) is a set of CC winning committees.*

Decisiveness of the Rules on the Paths. Consider the next example.

Example 6. *In Examples 2 and 3 we have already seen that p-HarmonicBorda and ℓ_p-Borda rules can provide results that refine the CC rule; the same happens for the k-Borda rule. Consider an election with $m \geq 2k$ candidates and two votes: $c_1 \succ c_2 \succ \cdots \succ c_m$ and its reverse. For the target committee size $2k$, the $2k$-Borda rule outputs all committees of size $2k$. On the other hand, each ℓ_p-Borda rule (with $p > 1$) and each p-HarmonicBorda rule (with $p > 0$) output a unique winning committee $\{c_1, \dots, c_k, c_m, \dots, c_{m-k+1}\}$.*

Indeed, our new rules can quite substantially improve the decisiveness of k-Borda and CC. We have generated $30,000$ elections with 10 candidates and 10 voters each (using the Impartial Culture assumption, i.e., we chose each voter's preference order uniformly at random) and computed all winning committees of size $k = 4$. Table 1 shows the results that we have obtained (we used brute-force search to compute all winning committees). Indeed, it seems that (the limiting variants of) the ℓ_p-Borda and p-HarmonicBorda rules can be used as natural, more decisive, variants of k-Borda and CC.

4 Experimental Results

In this section the goal is to illustrate our three paths between k-Borda and CC using the recent visual approach of Elkind et al. [2017a] (who computed graphical representations of election results for a number of rules and distributions of the candidates and voters in a two-dimensional space).

rule	avg. num. of win. committees	max. num. of win. committees
k-Borda	1.206	10
$1/10$-HarmonicBorda	1.00003	2
$\ell_{1.01}$-Borda	1.00	1
CC	2.119	28
5-HarmonicBorda	1.00016	2
ℓ_{50}-Borda	1.00007	2

Table 1: Average number of winning committees and the highest number of winning committees observed; under Impartial Culture assumption with $n = m = 10$ and $k = 4$.

In the 2D Euclidean model of preferences, voters and candidates are represented as points on a plane (often referred to as the ideal points of the respective voters/candidates). Each voter forms a preference order by sorting the candidates with respect to the distance of their ideal points from his or hers.

In our experiments—following Elkind et al. [2017a]—we consider elections with $m = 200$ candidates, $n = 200$ voters, and with committee size $k = 20$. These numbers are chosen since they are big enough to be quite representative and robust, but small enough to make experimentation feasible. All our elections are generated using the 2D Euclidean model, where both the candidates' and voters' ideal points are distributed uniformly on a $[-3, 3] \times [-3, 3]$ square. Elkind et al. [2017a] have considered other distributions and other election sizes, showing that the graphical results are robust to such variations (unless we consider, e.g., very small committees).

Computing Winning Committees. Since our rules are NP-hard, we used the following heuristic based on simulated annealing. We begin by sampling a random committee S_0. Then, in each iteration i, we take the committee S_{i-1} and form a temporary committee S'_i by randomly replacing one member in S_{i-1}. If the score of S'_i is greater than that of S_{i-1}, then we set S_i to be S'_i. Otherwise, we draw a random number between 0 and 1; if it is below pq^i (where p and q are two parameters, we used $p = 0.2$ and $q = 0.999$), then we set $S_i = S'_i$; otherwise, we set $S_i = S_{i-1}$. We execute 2000 iterations and output the highest-scoring committee encountered.

We compared the results obtained from using our heuristic to the optimal ones (for some of our OWA-based rules) and they turned out to be very close (both in terms of the approximation ratio and similarity of the histograms; see below).

The Experiment. We have generated $5,000$ elections for each of our rules and computed their results using our heuristic (we considered t-Borda rules for $t \in \{1, 2, \dots, 20\}$, p-HarmonicBorda for $p \in \{0.1, 0.2, \dots, 5.0\}$, and ℓ_p-Borda rules for $p \in \{1, 2, \dots, 50\}$). For each of the rules, we have computed a histogram showing how frequently candidates in various areas of the $[-3, 3] \times [-3, 3]$ square are selected. Specifically, we partitioned the square into 120×120 small squares (called cells) and computed how many winners from the generated elections fall into each. We present (some of) the resulting histograms in Figures 1–3 (the darker a given cell, the more winners fall there; Elkind et al. give more details on converting numerical values to colors). For comparison, Figure 4 contains histograms for k-Borda, Bloc, and CC.

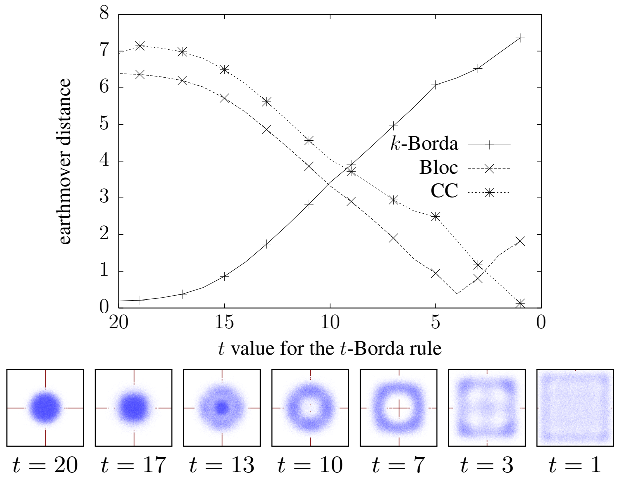

Figure 1: Results for the t-Borda rules ($k = 20$).

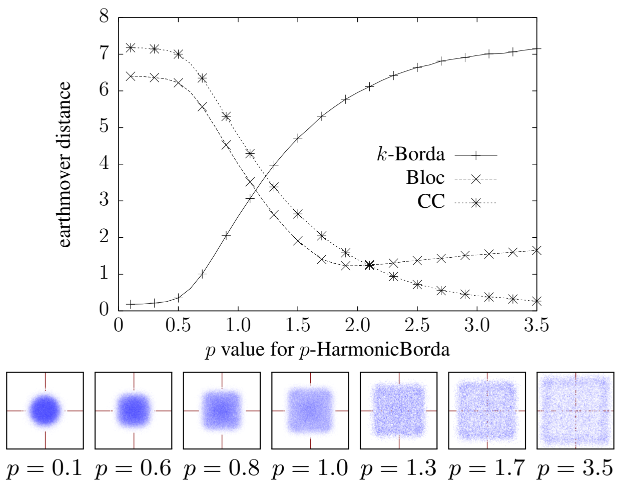

Figure 2: Results for the p-HarmonicBorda rules.

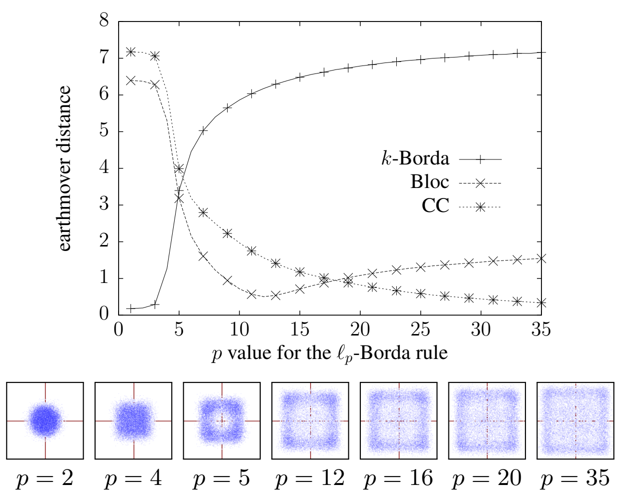

Figure 3: Results for the ℓ_p-Borda rules.

Finally, for each of the generated histograms, we computed its earth-mover distance from the histograms for k-Borda, Bloc, and CC. The plots showing how these distances change depending on our rules' parameters are shown in Figures 1–3.

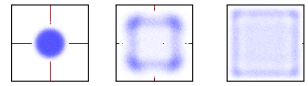

Figure 4: Histograms for k-Borda, Bloc, and CC.

Earth-Mover Distance. The earth-mover distance (EMD) is a widely-used measure of similarity between pictures [Peleg *et al.*, 1989]. Figuratively speaking, it views each cell in a histogram as containing some grains of sand (in proportion to the value of the cell), and the distance between two histograms is the minimum sum of Euclidean distances that all the grains of sand need to travel to transform one histogram into the other. To compute earth-mover distances, we used its standard formulation as an integer linear program (ILP) and used an ILP solver. The smaller the earth-mover distance between two histograms, the more similar these histograms are. The EMDs computed for our histograms are all normalized in the same way, so the results are comparable between rules. To assess which values of the distance are "small" and which are "large," one may compare presented example histograms.

Analysis. Perhaps somewhat surprisingly, the graphical results that we have obtained for the t-Borda and ℓ_p-Borda rules progress from k-Borda to CC through histograms very similar to those for Bloc (however, there is a big difference between Bloc and these rules; for each single election, Bloc typically chooses candidates concentrated in some parts of the square—leading to the histogram from Figure 4 only in the aggregate—whereas winners under t-Borda and ℓ_p-Borda rules are distributed very similarly to what their histograms suggest; the same good property holds for the p-HarmonicBorda rules). Further, for ℓ_p-Borda we observe a rather abrupt transition from histograms very similar to those for k-Borda (for $p \in \{1, \dots 4\}$) to those similar to Bloc (for $p \in \{5, \dots, 17\}$). The further transition, however, is smooth.

On the other hand, the histograms for p-HarmonicBorda rules do not pass though areas of similarity to Bloc (even though the histograms are more similar to those for CC than those for Bloc only for $p > 2$). Indeed, while the histograms for t-Borda and ℓ_p-Borda have areas of "lower density in the center," this never happens for p-HarmonicBorda.

5 Conclusions

We analyzed three families of committee scoring rules that form continuous paths between the k-Borda and CC rules. We have shown that the rules on these paths can lead to more decisive variants of k-Borda and CC. Further, these rules may be useful for tasks that need some compromise between excellence and diversity, provided by k-Borda and CC, respectively. Quite surprisingly, we have found that two of our paths seem to pass through rules with unexpected features.

For future research, we believe it would be interesting to consider paths between other multiwinner rules (e.g., between approval-based rules).

References

[Aziz *et al.*, 2017] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.

[Barberà and Coelho, 2008] S. Barberà and D. Coelho. How to choose a non-controversial list with k names. *Social Choice and Welfare*, 31(1):79–96, 2008.

[Betzler *et al.*, 2013] N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research*, 47:475–519, 2013.

[Chamberlin and Courant, 1983] B. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733, 1983.

[Debord, 1992] B. Debord. An axiomatic characterization of Borda's k-choice function. *Social Choice and Welfare*, 9(4):337–343, 1992.

[Elkind and Ismaili, 2015] E. Elkind and A. Ismaili. OWA-based extensions of the Chamberlin-Courant rule. In *Proceedings of ADT-15*, pages 486–502, September 2015.

[Elkind *et al.*, 2017a] E. Elkind, P. Faliszewski, J. Laslier, P. Skowron, A. Slinko, and N. Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of AAAI-17*, 2017.

[Elkind *et al.*, 2017b] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48(3):599–632, 2017.

[Faliszewski *et al.*, 2016a] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Committee scoring rules: Axiomatic classification and hierarchy. In *Proceedings of IJCAI-16*, pages 250–256, 2016.

[Faliszewski *et al.*, 2016b] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner analogues of the plurality rule: Axiomatic and algorithmic views. In *Proceedings of AAAI-16*, pages 482–488, 2016.

[Faliszewski *et al.*, 2016c] P. Faliszewski, A. Slinko, K. Stahl, and N. Talmon. Achieving fully proportional representation by clustering voters. In *Proceedings of AAMAS-16*, pages 296–304, 2016.

[Goldsmith *et al.*, 2014] J. Goldsmith, J. Lang, N. Mattei, and P. Perny. Voting with rank dependent scoring rules. In *Proceedings of AAAI-14*, pages 698–704, July 2014.

[Gordon and Truchon, 2008] S. Gordon and M. Truchon. Social choice, optimal inference and figure skating. *Social Choice and Welfare*, 30(2):265–284, 2008.

[Kilgour and Marschall, 2012] M. Kilgour and E. Marschall. Approval balloting for fixed-size committees. In M. Machover D. Felsenthal, editor, *Electoral Systems: Paradoxes, Assumptions, and Procedures*. Springer, 2012.

[Lang and Skowron, 2016] J. Lang and P. Skowron. Multi-attribute proportional representation. In *Proceedings of AAAI-16*, pages 530–536, February 2016.

[Lu and Boutilier, 2011] T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of IJCAI-11*, pages 280–286, 2011.

[Peleg *et al.*, 1989] S. Peleg, M. Werman, and H. Rom. A unified approach to the change of resolution: Space and gray-level. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):739–742, 1989.

[Peters and Elkind, 2016] D. Peters and E. Elkind. Preferences single-peaked on nice trees. In *Proceedings of AAAI-16*, pages 594–600, 2016.

[Procaccia *et al.*, 2008] A. Procaccia, J. Rosenschein, and A. Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30(3):353–362, 2008.

[Procaccia *et al.*, 2012] A. Procaccia, S. Reddi, and N. Shah. A maximum likelihood approach for selecting sets of alternatives. In *UAI*, pages 695–704, 2012.

[Ratliff and Saari, 2014] T. Ratliff and D. Saari. Complexities of electing diverse committees. *Social Choice and Welfare*, 43(1):55–71, 2014.

[Skowron *et al.*, 2015a] P. Skowron, P. Faliszewski, and A. Slinko. Achieving fully proportional representation: Approximability result. *Artificial Intelligence*, 222:67–103, 2015.

[Skowron *et al.*, 2015b] P. Skowron, L. Yu, P. Faliszewski, and E. Elkind. The complexity of fully proportional representation for single-crossing electorates. *Theoretical Computer Science*, 569:43–57, 2015.

[Skowron *et al.*, 2016a] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216, 2016.

[Skowron *et al.*, 2016b] P. Skowron, P. Faliszewski, and A. Slinko. Axiomatic characterization of committee scoring rules. Technical Report arXiv:1604.01529 [cs.GT], April 2016.

[Thiele, 1895] T. Thiele. Om flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, pages 415–441. 1895.

[Wen and Zhang, 2006] J. Wen and Z. Zhang. Vandermonde-type determinants and inequalities. *Applied Mathematics E-Notes*, 6:211–218, 2006.

[Yu *et al.*, 2013] L. Yu, C. Hau, and E. Elkind. Multiwinner elections under preferences that are single-peaked on a tree. In *Proceedings of IJCAI-13*, volume 13, pages 425–431, 2013.