

Restricted Preference Domains in Social Choice: Two Perspectives

Edith Elkind

University of Oxford, UK

Abstract. Preference aggregation is a challenging task: Arrow’s famous impossibility theorem [1] tells us that there is no perfect voting rule. One of the best-known ways to circumvent this difficulty is to assume that voters’ preferences satisfy a structural constraint, such as, e.g, being single-peaked. Indeed, under this assumption many impossibility results in social choice disappear. Restricted preference domains also play an important role in computational social choice: for instance, there are voting rules that are NP-hard to compute in general, but admit efficient winner determination algorithms when voters’ preferences belong to a restricted domain. However, restricted domains that have nice social choice-theoretic properties are not necessarily attractive from an algorithmic perspective, and vice versa. In this note, we will discuss some domain restrictions that have proved to be useful from a computational perspective, and compare the use of restricted domains in computational and classic social choice theory.

1 Introduction

A family of three—Alice, Bob, and their daughter Claire—would like to decide what to do on a hot Sunday morning. They are choosing among a bike ride, a trip to the pool, and a visit to the farmers’ market, and they only have time for one activity. Bob prefers the bike ride to the pool, and the pool to the farmers’ market. Alice prefers the farmers’ market to the bike ride, and the bike ride to the pool. Finally, Claire prefers the pool to the farmers’ market, and the market to the bike ride.

Aggregating these preferences into a decision that will be accepted by all family members is not easy: whichever option is chosen, there will be another option that would be preferred by a majority of the family members to the original choice. Moreover, the collective preferences exhibit a cyclic structure: a majority (Bob and Alice) prefer the bike ride to the pool, a majority (Bob and Claire) prefer the pool to the farmers’ market, and a majority (Alice and Claire) prefer the farmers’ market to the bike ride.

This simple example illustrates a fundamental difficulty in preference aggregation, and underpins a classic result in social choice: Arrow’s impossibility theorem [1], which tells us that there is no universally acceptable way of aggregating individual preferences into a collective decision. Indeed, social choice theory is unable to help Alice, Bob and Claire in this scenario unless they are willing to use randomization.

Suppose, on the other hand, that Alice, Bob and Claire decide to stop fighting, stay home, turn on the air conditioner and watch a movie. There is still a decision to be made, namely, how to set the room temperature. Alice’s preferred option is +21, Bob

favors +23, and Claire is happy with +25. It is then safe to assume that Alice also prefers +23 to +25, and Claire prefers +23 to +21. Hence, if we select +23, a majority of the voters would prefer that choice to +21, and a majority of the voters would prefer it to +25. Thus, it makes sense for them to agree on +23.

Why was it easier for Alice, Bob, and Claire to decide on the room temperature, compared to selecting a joint activity? Were they lucky, or was there something fundamental about the structure of the problem that ensured the existence of an acceptable alternative? Perhaps there is always a good solution for a family of three, but a bigger family would struggle to identify a majority-supported alternative? It turns out that there is a fundamental difference between the two examples, which was identified by Black back in 1948 [3]: temperature is, fundamentally, a one-dimensional concept, so anyone whose top choice is x would prefer $x + 1$ to $x + 2$, and $x - 1$ to $x - 2$. As demonstrated by Black, this rules out the possibility that collective preferences are cyclic, for any number of voters. Moreover, there always exist a choice that is preferred to any other choice by at least half of the voters.

The preferences studied by Black are formally known as *single-peaked preferences*: the space of alternatives can be viewed as a line, and when we graph each voter's intensity of preferences, the resulting curve has a single peak. There are other preference restrictions with similar properties, though the single-peaked domain is perhaps the most famous one. Restricted preference domains that circumvent Arrow's impossibility result and other similar impossibility theorems have long been a subject of study in social choice theory.

Now, in the examples considered so far, computation was not an issue: with just three voters and three alternatives one can easily implement any reasonable choice mechanism. However, as the number of voters and alternatives increases, algorithmic considerations become important. Consider, for instance, a student computer club that needs to decide on the slate of seminar speakers. There is a limited number of slots, and each club member has ranked all potential speakers; the goal is then to select a subset of speakers so that there is something in the seminar program for every student. One can formalize this goal and define a voting rule that selects an optimal set of speakers according to it; this rule is known as the Chamberlin–Courant rule [4]. However, winner determination under this rule is known to be NP-hard [9]. Inspired by Black's results we can then ask: does this hardness result survive if voters' preferences are single-peaked? It turns out that the answer is no: for single-peaked preferences, the output of the Chamberlin–Courant rule and other similar rules can be computed in polynomial time using a dynamic programming approach [2].

Inspired by this and similar observations, the computational social choice community began to systematically investigate whether existing hardness results in computational social choice still hold when voters' preferences belong to a restricted domain, with a focus on domains defined in the social choice literature. The answer to this question turned out to be more nuanced than one might expect: while classic domain restrictions, such as single-peaked or single-crossing preferences, turn out to be incredibly useful from an algorithmic perspective, there are other domains that have some appealing social choice properties, but do not enable efficient algorithms for important social choice problems, as well as domains that do not rule out cyclic preferences, but

are nevertheless algorithmically useful. In this talk, we will consider several such examples, including, in particular, preferences single-peaked on trees [5, 13, 11], preferences single-peaked on a circle [12], and single-peaked preferences in the approval setting [6].

2 Preliminaries and Notation

For every positive integer n , set $[n] = \{1, \dots, n\}$. Let C be a finite set of *alternatives*, or *candidates*, and let $m = |C|$. A *linear order* over C is a binary relation over C that is complete, transitive and antisymmetric. Given a linear order v over C , we denote the top alternative in v by $\text{top}(v)$.

Definition 1. A profile $P = (v_1, \dots, v_n)$ over a set of alternatives C is a list of linear orders over C . We associate $P = (v_1, \dots, v_n)$ with a set of voters $N = [n]$; the order v_i is called the vote of voter i . For convenience, we write $a \succ_i b$ whenever $(a, b) \in v_i$, i.e., when voter i strictly prefers a to b .

Given a profile C over A , we define its *weak majority relation* \succeq_{maj} as a binary relation over C such that

$$a \succeq_{\text{maj}} b \iff |\{i \in N : a \succ_i b\}| \geq |\{i \in N : b \succ_i a\}|.$$

We write $a \succ_{\text{maj}} b$ if $a \succeq_{\text{maj}} b$, but not $b \succeq_{\text{maj}} a$. Alternative a is a *weak Condorcet winner* if $a \succeq_{\text{maj}} b$ for all $b \in C$; it is a *Condorcet winner* if $a \succ_{\text{maj}} b$ for all $b \in C$.

Single-Peaked Preferences Let \triangleleft be a linear order over the set of alternatives C . A vote v over C is *single-peaked with respect to* \triangleleft if for every pair of candidates $a, b \in C$ with $\text{top}(v) \triangleleft b \triangleleft a$ or $a \triangleleft b \triangleleft \text{top}(v)$ it holds that v ranks b above a . A profile P over C is *single-peaked with respect to* \triangleleft if every vote in P is single-peaked with respect to \triangleleft ; P is *single-peaked* if there exists a linear order \triangleleft over C such that P is single-peaked with respect to \triangleleft . We refer to any such order \triangleleft as an *axis* for P .

It is known that if voters' preferences are single-peaked then the weak majority relation is transitive; in particular, this means that single-peaked profiles always have weak Condorcet winners.

The Chamberlin–Courant rule We will now describe a family of voting rules that will be used to illustrate the algorithmic properties of various restricted preference domains considered in this note.

Rules in this family take a candidate set C , a profile P over this set and a target committee size k as an input, and output a subset of candidates (committee) of size k . Given a candidate set C , $|C| = m$, every vector $\mathbf{s} = (s_1, \dots, s_m)$ of non-negative integers with $0 = s_1 \leq \dots \leq s_m$ defines a *positional misrepresentation function* $\mu_{\mathbf{s}} : P \times 2^C \rightarrow \mathbb{Z}$ as follows: $\mu_{\mathbf{s}}(v, C') = s_i$ if v ranks her most preferred candidate in C' in position i . The (utilitarian version of the) Chamberlin–Courant rule outputs some committee C' of size k that minimizes the quantity $\sum_{v \in P} \mu_{\mathbf{s}}(v, C')$ (which we call the *s-score* of C') over all size- k subsets of C . The misrepresentation function associated with the vector $\mathbf{s} = (0, 1, \dots, m-1)$ is known as the *Borda misrepresentation function*.

Finding a winning committee under the Chamberlin–Courant rule is known to be NP-hard, even for the Borda misrepresentation function [9]. In contrast, this problem can be solved in polynomial time for an arbitrary misrepresentation function if the input profile is single-peaked [2]. The algorithm of Betzler et al. [2] proceeds by dynamic programming: it considers the single-peaked ordering \triangleleft of the set of alternatives C , and, for every $m' \in [m]$ and every $k' \in [k]$, identifies the minimum misrepresentation among all committees of size k' that are contained in the m' -prefix of \triangleleft and contain the m' -th candidate. This approach works, because it is easy to measure to benefit accomplished by adding candidate m' to a committee whose rightmost member is r , $r < m'$.

In what follows, we will consider several extensions of the concept of single-peaked preferences, and analyze the complexity of computing the Chamberlin–Courant winners for preferences belonging to these domains.

3 Preferences Single-Peaked on a Tree

. The domain of preferences single-peaked on a tree was introduced by Demange [5]. Consider a profile P over a set of candidates C and a tree T with vertex set C . The profile P is said to be *single-peaked on T* if it is single-peaked on every path in T . Equivalently, for every vote $v \in P$ and every $s \in [m]$ it holds that the set of top s candidates in v is a connected subset of vertices of T . Clearly, a single-peaked profile is single-peaked on a tree (namely, the path associated with the underlying order \triangleleft), but the converse is not true.

Demange [5] has shown that if a profile is single-peaked on some tree then it has a weak Condorcet winner. However, its weak majority relation need not be transitive. In fact, it is not hard to see that for every tree that is not a path there exists a profile that is single-peaked on that tree, but whose majority relation is not transitive. Indeed, if a tree T is not a path, it has a vertex a with three neighbors b , c , and d . If a voter ranks a first, followed by b , c , and d (in any order), followed by all other candidates, then her preferences are single-peaked on T . Thus, we can construct a profile single-peaked on T where a majority of the voters prefers b to c , a majority prefers c to d and a majority prefers d to b . Hence, from a social choice perspective profiles that are single-peaked on trees have some desirable properties; however, there is a clear boundary between paths and other kinds of trees.

However, from a computational social choice perspective, the picture is very different. The dynamic programming algorithm of Betzler et al. [2] can be extended to trees, but its running time then becomes exponential in the number of leaves [13]: intuitively, whenever we try to compute an optimal committee of size k' contained in a subtree rooted at a , we have to consider the subtrees rooted at children of a , and, for each child, guess the topmost committee member in its subtree. Indeed, it can be shown that computing Chamberlin–Courant winners remains NP-hard when preferences are single-peaked on a tree, for a large class of misrepresentation functions. The hardness result holds even if the underlying tree is a star or has maximum degree 3. Thus, computing the Chamberlin–Courant winners is easy for long skinny trees, but not for bushy trees. However, this hardness result does not apply to the Borda misrepresentation function; indeed, for this misrepresentation function computing the Chamberlin–Courant winners

is easy if the profile is single-peaked on a star, or, more generally, on a tree with a small number of internal vertices [11].

Thus, for a social choice theorist all trees that are not stars have the same properties, whereas a computational social choice theorist distinguishes between ‘nice’ trees and arbitrary trees.

4 Preferences Single-Peaked on a Circle

Peters and Lackner [12] extend the notion of single-peaked preferences from paths to circles: under their definition, a profile P over a set of candidates C is single-peaked on a circle with vertex set C if for each voter $v \in P$ the circle C could be cut so that v is single-peaked on the resulting path. We immediately note that the preferences of Alice, Bob and Claire over morning activities are single-peaked on a circle, and hence we cannot expect profiles that are single-peaked on a circle to always have weak Condorcet winners or to guarantee the transitivity of weak majority preferences. This explains why this model has not been considered in the social choice literature.

On the other hand, preferences single-peaked on a circle turn out to be simple enough to be algorithmically useful: Peters and Lackner establish that for such preferences one can compute a winning committee under the Chamberlin–Courant rule in polynomial time. However, it is not clear how to use Betzler et al.’s dynamic programming algorithm for this purpose: instead, Peters and Lackner build on the techniques developed by Peters [10] to argue that the problem of finding Chamberlin–Courant winners in this scenario can be encoded as an integer linear program whose constraint matrix is totally unimodular.

5 Single-Peaked Preferences in the Approval Setting

So far, we assumed that voters’ preferences are expressed by linear orders over alternatives. Instead, one can consider a simpler scenario: each voter approves a subset of alternatives and disapproves the remaining alternatives. Such preferences are known as approval preferences. The setting of approval preferences is considered to be fairly simple from a social choice perspective: if the goal is to select a single winner, the rule that selects an alternative with the highest number of approvals has a number of attractive normative properties (see, e.g., [8] and the references therein). Thus, domain restrictions in the context of approval preferences have not been considered in the social choice literature. However, voting with approval preferences still presents interesting algorithmic challenges; in particular, a natural adaptation of the Chamberlin–Courant rule for this setting remains NP-hard. On the other hand, for single-peaked approval profiles (where alternatives can be placed on a line so that each voter approves a contiguous segment of this line) Chamberlin–Courant winners can be computed efficiently by dynamic programming [6].

6 Conclusion

We have seen three examples of restricted domains that would be viewed very differently by a social choice theorist and an algorithms researcher: in the first case, the algorithmic approach offers a more fine-grained view of the domain, and in the second and third case, a domain restriction that is seen as not particularly useful from a social choice perspective (either because it does not eliminate violations of the Condorcet principle or because the underlying domain is considered to be simple enough to start with) turns out to be useful from an algorithmic perspective. In this note, we focused on the Chamberlin–Courant rule; one can consider other computationally challenging voting rules, in which case the picture becomes even more complicated. For more details, we refer the reader to the recent survey [7].

Acknowledgments This work was supported by the European Research Council (ERC) under grant number 639945 (ACCORD).

References

1. Kenneth J. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 1951.
2. Nadja Betzler, Arkadii Slinko, and Johannes Uhlmann. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research*, 47(1):475–519, 2013.
3. Duncan Black. On the rationale of group decision-making. *The Journal of Political Economy*, 56(1):23–34, 1948.
4. John R. Chamberlin and Paul N. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733, 1983.
5. Gabrielle Demange. Single-peaked orders on a tree. *Mathematical Social Sciences*, 3(4):389–396, 1982.
6. Edith Elkind and Martin Lackner. Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2019–2025, 2015.
7. Edith Elkind, Martin Lackner, and Dominik Peters. Structured preferences. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 10, pages 187–207. AI Access, 2017.
8. Ulle Endriss. Sincerity and manipulation under approval voting. *Theory and Decision*, 74(3):335–355, 2013.
9. Tyler Lu and Craig Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 280–286, 2011.
10. Dominik Peters. Single-peakedness and total unimodularity: New polynomial-time algorithms for multi-winner elections. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1169–1176, 2016.
11. Dominik Peters and Edith Elkind. Preferences single-peaked on nice trees. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 594–600, 2016.
12. Dominik Peters and Martin Lackner. Preferences single-peaked on a circle. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 649–655, 2017.
13. Lan Yu, Hau Chan, and Edith Elkind. Multiwinner elections under preferences that are single-peaked on a tree. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 425–431, 2013.