Voting with Limited Energy: A Study of Plurality and Borda

Zoi Terzopoulou
GATE, University of Lyon - Saint-Etienne
France

ABSTRACT

An abundance of problems rely on voting, ranging from standard political elections and committee decisions to coordinated efforts of multiagent systems. A common and prevalent, yet often underestimated, element of these situations is the substantial effort required by the voters to examine all alternatives involved and to form complete preferences. How does limited energy affect collective decision making? This is the question we address, enriching the classical framework of voting by incorporating two new parameters: the energy limits of the voters, as well as the order in which the alternatives are presented to them. We focus on two popular voting rules: Plurality and Borda. We conduct an extensive social welfare analysis with both analytical and experimental tools, and we also study the strategic incentives that arise in this setting.

KEYWORDS

Social Choice; Voting; Limited Energy; Incomplete Information

ACM Reference Format:
Zoi Terzopoulou. 2023. Voting with Limited Energy: A Study of Plurality and Borda. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

Voting is a tool that plays an increasingly important role in the fields of artificial intelligence and multiagent systems: From aggregating the preferences of artificial agents in order to decide about a joint action to employing online platforms in order to find a consensus amongst multiple users, methods for collective decision making that were in the past solely studied by economists are establishing their value within ample modern applications [9]. Accurately modelling how voters form their opinions is an essential factor for the success of such endeavours, which either rely on human behaviour directly, or are based on simulations of human behaviour through AI.

Yet, research on *computational social choice* so far severely ignores decision-making patterns that are exhibited by real people, and constraints itself in stringent assumptions that would rarely be relevant in practice [3]: A vast majority of formal frameworks assume that agents in a group will report preferences over the full set of alternatives in a given voting setting, no matter how large that set may be. This ideal scenario is arguably difficult to obtain in reality: people with limited energy, finite time, or bounded attention cannot be expected to go through numerous pages of options before reporting their preferences regarding a meeting date in the *doodle* platform, a neighbourhood project in a *participatory budgeting* experience, a hotel in a travel *website*, a promising candidate in an *election*, or a paper in a conference *bidding* phase [5].

When agents with limited energy report an incomplete part of their intrinsic preferences, do the collective outcomes bear negative consequences? If so, are there ways in which we—as mechanism designers—can alleviate them? These are some of the questions that this paper targets. Example 1 provides further motivation.

Example 1. Consider a group of five friends about to order dinner from a delivery app. The app includes three options, in the order x, y, z. Two patient friends examine all options in detail and form their preferences: they agree that z is the best (followed by y and then by x). Being more impatient, the remaining three friends stop looking at options after x and y; they know then that they rank the former above the latter. The preferences and the votes are depicted in Figure 1, with the votes appearing within shaded boxes.[1] Suppose that after reporting their preferences, the friends select the option that is found to be the most desirable one by most of them: option x. However, had all the friends spent the energy examining all options, they would realise that z is actually everyone's favourite—instead, they will all have to accept a worse outcome.

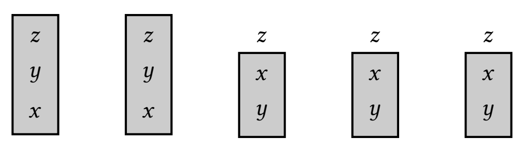

Figure 1: A negative effect of limited energy in voting.

Despite its simplicity, Example 1 brings to the surface all parameters that are relevant for our study: How do the intrinsic preferences of the voters relate to each other, and what kind of energy limits prohibit voters from accessing them in full?[2] Do all voters examine their options in the same order, and what method (i.e., *voting rule*) does the group use to reach a collective decision? We will focus on two famous voting rules, which operate on different aspects of the voters' preferences and may produce different outcomes: The *Borda* rule and the *Plurality* rule (note that the latter was implicitly applied in Example 1) [35]. Plurality relies on little information, looking at the most preferred alternative of each voter; Borda on the other hand needs more information, equally distributing points to the alternatives based on their position in a voter's preference. Contrasting these two voting methods will give us precious insights.[3]

[1]We will follow this drawing convention throughout the paper.
[2]In Example 1 there is a high degree of concordance between the preferences of the voters and between their energy limits. Even then, we will see that choosing the order in which the alternatives should be presented to the voters is not straightforward.
[3]Comparing Plurality and Borda is standard in social choice. See e.g., [4, 15].

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

To measure the effects of limited energy on a group of voters, we will employ a notion of social welfare from the literature [14, 24]. Our venture is high-reaching for the practical relevance of voting theory within digital environments: If these effects are found to be negative and strong, then either the design of a voting platform needs to be revised, or the voters must be incentivised to devote more energy when reflecting on their preferences.

We must be careful though: Everything we have discussed so far hinges on a standard assumption of the social choice community—that the voters are *sincere* and simply express their preferences with their votes. But voters may also act strategically, reporting untruthful preferences towards an outcome that they consider better. Does limited energy offer any protection against such strategic behaviour? By enriching existing work on voting with incomplete information, we will later reveal an answer to this question.

Our Contribution. This paper proposes a novel *formal framework* for the study of voting under limited energy. The landscape we explore is critical for understanding voting under non-ideal conditions and for enhancing existing mechanisms. We analytically compare two voting rules that are central in the field of computational social choice. Our theoretical contributions regard issues of social welfare, as well as the voters' incentives for strategic manipulation. In addition, we provide a comprehensive experimental study, scrutinising a rich domain of voting contexts. To that end we have built a voting *app*, where simulation experiments can be conducted by the users easily, enabling them to directly confirm our results. The app can be found in https://github.com/zoi-ter/voting-with-limited-energy, together with the full version of this paper. The Appendix is also included in that version.

Related Work. Several streams of literature are related to our framework, the most pertinent of which is probably *preference elicitation*. Elicitation methods are used to gain information about the preferences of the voters when access to the complete preferences is costly, either for the voters themselves or for the mechanism designer. The parting point with our work is that preference elicitation techniques are commonly incremental and interactive, and aim at finding a suitable collective outcome with minimum queries [2, 16]. However, we (i) do not have the option to elicit more preference information from the voters than what their energy permits, (ii) do not benefit from asking fewer questions than what the voters are able to answer, and (iii) cannot modify our pre-selected process in accordance with the preferences we observe (the order in which the alternatives are presented to the voters—e.g., in a query about a restaurant in a delivery app—is fixed in advance and the voters report their preferences in a single shot). More broadly, *communication complexity* is not a concern in our setting.

Another related research theme is that of the *possible and necessary winners* [17], asking which alternatives may win under complete preferences when we only have incomplete information about those preferences. It is known that in the worst case, winners of many voting rules cannot be determined without a large amount of information being provided by the voters [7]. Lu and Boutilier [21] propose a method based on minmax regret to select a winner given incomplete preferences, but this kind of methods do not use additional information about the origin of incompleteness (which we have in our framework from the energy limits), and they consider all complete extensions of the reported preferences. Instead, we want to apply the voting rules *directly* on the incomplete preferences. Technically, our setting could be studied as a subcase of the general possible winners problem; conceptually, it is easier to explain it to the voters (only requiring the explanation of a simple voting rule rather than of probabilistic calculations), and is also easier to trust it (because it does not use any assumptions about the distribution of the hidden preferences but only counts on the revealed information).

There is a progressively large literature on voting with incomplete preferences, from which we adopt some notation and terminology [6, 25, 33]. To the best of our knowledge, none of these works is concerned with the problem of limited energy, thus our targets are different. For instance, in a recent article, Ayadi et al. [1] define approximations for voting rules based on information only about the top-k alternatives from a voter's preference—asking for a top subset of someone's preference is not sensible in our setting, since it would mean that voters are always aware of these specific alternatives even if they are not presented with them.

Distortion is also a related concept that has attracted the interest of social choice researchers [27]. It commonly regards the effects that occur when we apply a voting rule to types of preferences that are not as refined as the intrinsic ones of the voters (e.g., when we use the ordinal information from the voters' cardinal utilities). Our paper is inspired by work on distortion, but for all we know similarities between the two frameworks remain only in spirit.

There is a plethora of evidence by social scientists and psychologists about the bounds in people's decision making, limited energy being one of them [12, 30]. Most germane to our paper is the cognitive heuristic of *satisficing*, which entails searching through the available alternatives until an acceptability threshold of the voter is met [32]. What we will later define as energy limits can be thought to emanate from satisficing.

Paper Overview. Section 2 presents our framework, together with notation and terminology. Section 3 includes our analytical, worst-case results on social welfare, and Section 4 elaborates on this topic by simulation experiments. Section 5 continues with the investigation of voters' strategic behaviour in our setting, and Section 6 concludes with a summary of our work and open questions.

2 THE MODEL

The basic ingredients of a voting problem are a group of *voters* $N = \{1, \ldots, n\}$ and a set of *alternatives* $A = \{x_1, \ldots, x_m\}$. Every voter $i \in N$ is associated with an *intrinsic* preference ranking $\succ_i$ (i.e., a linear order) over *all* alternatives in A. By $x \succ_i y$ we mean that voter i prefers $x \in A$ over $y \in A$, conditionally that she is aware of both alternatives. If $x \succ_i x_1 \succ_i \cdots \succ_i x_{k-1} \succ_i y \succ_i \cdots \succ_i x_m$ for some $x_1, ..., x_{k-1} \in A$, we say that x is on the first level of $\succ_i$, x_1 on the second level, etc., and that x is k levels *above* y. A *profile* is a vector with the preferences of all voters, denoted by $\boldsymbol{\succ} = (\succ_1, \ldots, \succ_n)$. We also write $\boldsymbol{\succ}_{-i}$ for the profile of all voters' preferences except for i's. Given a subset of alternatives $S \subset A$, we write $\succ\restriction_S$ for the restriction of $\succ$ to S. Concretely, the following holds for all alternatives $x, y \in A$:

$$x \succ\restriction_S y \text{ if and only if } x \succ y \text{ and } x, y \in S$$

Then, $\succ\restriction_S = (\succ_1\restriction_S, \ldots, \succ_n\restriction_S)$ is the restriction of $\succ$ to S.

Energy Limits. Every voter $i \in N$ is equipped with an *energy limit* $e_i \in \{2, \ldots, m\}$, i.e., the number of alternatives she is capable to consider in the given voting problem.[4] An *energy function* $e : N \to \{2, \ldots, m\}$ captures the energy limit of every voter. We call an energy function e *non-trivial* if $e_i \neq m$ for all $i \in N$. Denoting by $\Pi(A)$ the set of all permutations of A, we also define a *show order* function $o : N \to \Pi(A)$, where for each voter i, $o_i \in \Pi(A)$ indicates the order in which the alternatives are presented to her. In a voting problem, every voter i only examines the first e_i alternatives from o_i; we denote this set by $A(e_i, o_i) \subseteq A$, and the induced (partial) preference ranking by $\succ_i^{e,o} = \succ_i\restriction_{A(e_i,o_i)}$.[5] This is the *reported preference* of voter i, viz., her *vote*.

Voting Rules. A *voting rule* is a function F that takes as input a profile of (possibly partial) preferences $\succ^{e,o} = (\succ_1^{e,o}, \ldots, \succ_n^{e,o})$ and outputs a winning alternative. In particular, a positional scoring rule F assigns a score $s_i^F(x)$ to every alternative x depending on its position in the reported preference $\succ^{e_i,o_i}$ of voter i, and selects as winner the alternative with the largest accumulated score for all voters.[6] Equivalently, for every voter i and alternative x, the score $s_i^F(x)$ can be induced from a scoring vector $s_i^F = (s_1^F, \ldots, s_{e_i}^F)$, meaning that s_1^F corresponds to the alternative in $A(e_i, o_i)$ appearing on the first level of the ranking $\succ_1^{e,o}$, s_2^F corresponds to the second level, etc. We make the typical assumption that ties are broken according to the lexicographic ordering of the alternatives.

$$F(\succ^{e_i,o_i}) \quad = \quad \operatorname*{argmax}_{x \in A} \sum_{i \in N} s_i^F(x)$$

Two classical positional scoring rules in voting theory are *Plurality* and *Borda*. Plurality gives score 1 to the first alternative in a preference ranking and score 0 to all other alternatives; on the other hand, the (symmetric) Borda score of an alternative x is the difference between the number of alternatives that are ranked below x and the number of alternatives that are ranked above x [35].

The scoring vectors corresponding to our two rules are defined as follows (where all alternatives in $A \setminus A(e_i, o_i)$, which are not examined by voter i, are assigned score 0):

$$s_i^{Plur} = (1, \overbrace{0, \ldots, 0}^{e_i - 1})$$

$$s_i^{Borda} = (e_i - 1, e_i - 3, \ldots, 0, \ldots, -e_i + 3, -e_i + 1)$$

Note that for the special case of total energy, (i.e., when $e_i = m$ for all $i \in N$), we obtain the standard definitions from voting theory.

Social Welfare. We first define the utility that a voter receives from a voting outcome, and subsequently measure the social welfare of the group as a whole. We denote voter i's *utility* (or *welfare*) when x is the winner by $w_i(x)$, and assume that it depends on the voter's intrinsic preference $\succ_i$ in an *equally distributed*, *linear* fashion:

$$w_i(x) = (m - 1) - |y \in A : y \succ_i x|$$

For instance, a voter gets 0 welfare if her intrinsically least preferred alternative wins, and gets $(m - 1)$ welfare if her intrinsically most preferred alternative wins. A debatable feature of this welfare notion is that it can also be seen as an alternative definition of the Borda scores (corresponding to an affine transformation of the scores defined above). Yet, we will stick to this notion from the literature [18, 23] and make an interesting observation: that the Borda rule, despite its apparent advantages, does not necessarily outperform Plurality. Then, the group's *social welfare* when alternative x wins is defined as $SW(x) = \sum_{i \in N} w_i(x)$.

We can compare the quality of the outcomes that our two different rules produce under limited energy in an *absolute* manner, by comparing the social welfare values obtained given a particular profile of preferences and a particular show order.[7] But we can also measure the collective effects of limited energy in a *relative* manner, i.e., by comparing the welfare produced by a rule under limited energy with the one that would be produced under full energy. To account for such relative worst-case effects for an energy function e, we define a notion called *Price of Limited Energy* (PLE):[8]

$$PLE(F, e) = \max_{\succ, o} \frac{SW(F(\succ))}{SW(F(\succ^{e,o}))}$$

The higher PLE is, the worse-off a group is if its members vote with limited energy in relation to voting with full energy.

3 THE EFFECTS OF LIMITED ENERGY: A THEORETICAL ANALYSIS

We are ready to answer the first important question of our paper: Is limited energy necessarily harmful for a group of voters? Proposition 1 is not difficult to prove: Limited energy necessarily damages a group's social welfare when the Borda rule is applied.

Proposition 1. *It holds that $SW(Borda(\succ)) \geq SW(Borda(\succ^{e,o}))$ for every profile $\succ$, energy function e, and show order function o.*

Proof. The key idea is that the social welfare corresponding to an alternative x is an affine transformation of the Borda score of x:

$$SW(x) = \sum_{i \in N} w_i(x) = \sum_{i \in N} \frac{s_i^{Borda}(x) + m - 1}{2}$$

This implies that $SW(Borda(\succ)) \geq SW(y)$ for all alternatives $y \in A$, including the alternative $y = Borda(\succ^{e,o})$. $\square$

On the contrary, when the Plurality rule is applied, limited energy can be deemed beneficial for a group of voters.

Proposition 2. *It holds that $SW(Plur(\succ^{e,o})) > SW(Plur(\succ))$ for some profile $\succ$, energy function e, and show order function o.*

Proof. Consider the profile $\succ$ below, where $m = 4$, $n = 5$, and $A(e_i, o_i) = \{x_2, x_3, x_4\}$ for all $i \in N$.

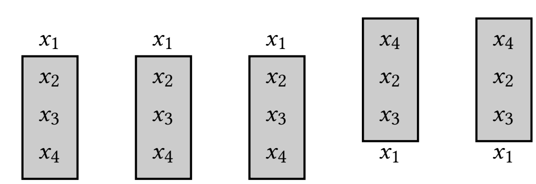

[4]For simplicity, we assume that every voter can consider at least two alternatives.
[5]The SOI data in the `preflib.org` website represent such strict incomplete orders.
[6]Kruger and Terzopoulou [20] give a definition of scoring rules for partial preferences.
[7]Absolute and normalised social welfare is presented in all figures of Section 4.
[8]Cf. the well-known game-theoretical notion of *Price of Anarchy* [19].

Then, the following proves the proposition: $SW(Plur(\succ^{e,o})) = SW(x_2) = 10 > 9 = SW(x_1) = SW(Plur(\succ))$. $\square$

We also find one further manner in which Plurality is superior to Borda: it guarantees that no matter the preferences and the energy of the voters, there always exist a show order such that the outcome that would be obtained if voters held full energy is preserved—this is not the case for the Borda rule (Propositions 3 and 4).

Proposition 3. *For every profile $\succ$ and energy function e, it holds that $Plur(\succ^{e,o}) = Plur(\succ)$ for some show order o.*

Proof. Choose a show order o such that $top_i \in A(e_i, o_i)$ for all i, where top_i is the most prefered alternative in $\succ_i$. $\square$

This is a significant advantage of Plurality if the intrinsic preferences of the voters are accessible—a rather unrealistic assumption. Still, these preferences may be approximated via AI, for example within frameworks of *preference learning* or *recommender systems*.

Proposition 4. *For some profile $\succ^{e,o}$ and energy funtion e, it holds that $Borda(\succ^{e,o}) \neq Borda(\succ)$ for all show orders o.*

Proof. Consider the profile $\succ = (\succ_1, \succ_2, \succ_3)$ where $x_1 \succ_1 x_2 \succ_1 x_3$, $x_2 \succ_2 x_3 \succ_2 x_1$, and $x_3 \succ_3 x_1 \succ_3 x_2$. The Borda winner (after tie-breaking) is $Borda(\succ) = x_1$. Suppose that $e_1 = 2$, while $e_2 = e_3 = 3$. There are three possibilities for the alternatives shown to voter 1, while the other two voters are shown all alternatives: First, $A(e_1, o_1) = \{x_1, x_2\}$, in which case $Borda(\succ^{e,o}) = x_3 \neq x_1$; second, $A(e_1, o_1) = \{x_2, x_3\}$, in which case $Borda(\succ^{e,o}) = x_2 \neq x_1$; third, $A(e_1, o_1) = \{x_1, x_3\}$, in which case $Borda(\succ^{e,o}) = x_3 \neq x_1$. $\square$

Although in some scenarios limited energy comes with social advantages for Plurality, Theorem 1 shows that in the worst case, limited energy is always damaging to social welfare, meaning that $PLE(Plur, e) \geq 1$ for every energy function e. In addition, Theorem 1 proves that even the slightest amount of missing information is enough to skew the collective outcome for both rules.

Theorem 1. *Let F be the Plurality rule or the Borda rule, and α some real number. Then:*

$$PLE(F, e) = \begin{cases} 1 & \text{if } e_i = m \text{ for all } i \in N \\ \alpha > 1 & \text{otherwise} \end{cases}$$

Proof. For an energy function where every voter has full energy, it is obvious that $PLE(F, e) = 1$. Suppose now that $e_i \neq m$ for some $i \in N$. Withoug loss of generality, take $e_1 \neq m$. Recalling that $e_i \geq 2$ for all $i \in N$, we construct a profile $\succ$ with $n = 3$ as shown in Figure 2, where $x_1 \notin A(e_1, o_1)$, $x_2 \in A(e_1, o_1)$, $x_1 \in A(e_2, o_2)$, and $x_2 \in A(e_3, o_3)$. As an illustration, we depict below a case where $e_2 = m - 1$ and $e_1 = e_3 = m$:

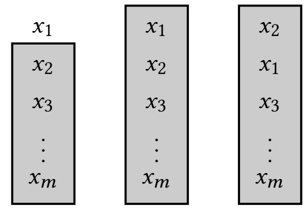

Figure 2: PLE increase under minimum lack of energy.

Clearly, $F(\succ) = x_1$ and $F(\succ^{e,o}) = x_2$ for F being either Plurality or Borda. Since $SW(x_1) > SW(x_2)$, we have that $PLE(F, e) > 1$. $\square$

Because our notion of social welfare is by definition maximised by the Borda rule in the special case of full energy, a question arises of whether the superiority of Borda is preserved under limited energy as well. Interestingly, we will next show that under limited energy, we may even observe opposite effects.

Proposition 5. *It holds that $SW(Plur(\succ^{e,o})) > SW(Borda(\succ^{e,o}))$ for some profile $\succ$, energy function e, and show order function o.*

Proof. Consider the profile $\succ$ below, where $m = 7$, $n = 3$, $A(e_i, o_i) = \{x_1, \ldots, x_6\}$ for $i \in \{1, 2\}$, and $A(e_3, o_3) = \{x_2, \ldots, x_7\}$.

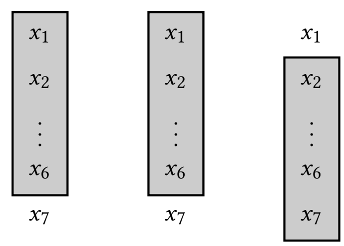

Then, we have that $SW(Plur(\succ^{e,o})) = SW(x_1) = 18 > 15 = SW(x_2) = SW(Borda(\succ^{e,o}))$. $\square$

In Section 4, the experiments will convince us that the profile in the proof of Proposition 5 does not constitute a rare example. This proof reveals a relevant intuition: Since Borda is sensitive to more information regarding voters' preferences, minor gaps in this information are able to alter the voting outcome—this is not true for Plurality outcomes, which can remain the same under more significant changes in the input profile. Yet, Plurality is inferior to Borda in the worst case: it produces a higher PLE (Theorem 2).

Lemma 1 is used in the proof of Theorem 2. Given a natural number $k \leq n$, we denote by $k.\min_{i \in N} e_i$ a set $S \in \{2, \ldots, m\}^k$ such that $|S| = k$ and $e_i \leq e_j$ for every $e_i \in S$ and $e_j \notin S$.

Lemma 1. *Let $\alpha = \lfloor \frac{n}{m-1} \rfloor + 1$ if $\lfloor \frac{n}{m-1} \rfloor \notin \mathbb{N}^+$ and $\alpha = \frac{n}{m-1}$ otherwise. The following holds for every non-trivial function e:*

$$PLE(Plur, e) \geq \frac{n(m-1)}{\sum_{\alpha.\min_{i \in N} e_i} e_i - 1}$$

Proof. We will prove that for any non-trivial energy function e we can construct a profile $\succ$ and a show order function o such that the required lower bound PLE for is realised.

The profile $\succ$ is built as follows: Let x be the lexicographically fist alternative in A. We choose an alternative $y \neq x$ and place it at the first position of every voter's ranking $\succ$, meaning that $SW(Plur(\succ)) = SW(y) = n(m-1)$. Then, we partition N into α subsets of voters, with as many subsets as possible having size $m-1$. In each of these subsets with size $m-1$, every alternative in $A \setminus \{y\}$ is ranked first in the preference $\succ_i^{e,o}$ of exactly one voter i; if one of these subsets has size $k \leq m-1$, then x and $k-1$ other alternatives from $A \setminus \{y\}$ are ranked first in the preference $\succ_i^{e,o}$ of exactly one voter i. Clearly, no alternative can have a strictly larger Plurality score than x in $\boldsymbol{\succ}^{e,o}$, so x is a Plurality winner under limited energy. To obtain the social welfare for x that appears in the denominator of the fraction in the statement, we ensure that whenever voter i does not rank x first in $\succ_i^{e,o}$, she will rank it last, and only the α voters with the smallest energy will rank x first. Thus, $SW(Plur(\boldsymbol{\succ}^{e,o})) = SW(x) = \sum_{\alpha.\,\min_{i\in N} e_i} e_i - 1$. Figure 3 exemplifies our construction with $m = 4$, $n = 7$, and $e_i = 2$ for all $i \in N$ (the remaining positions in the rankings are filled arbitrarily). □

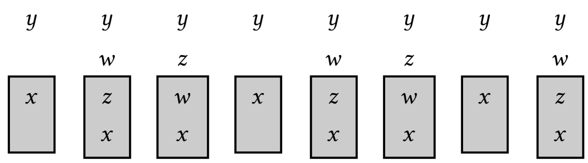

Figure 3: Profile construction for the PLE calculation of Plurality.

The lower bound for the PLE of the Plurality rule is not strict. See an example with $m = 4$, $n = 4$, and $e_i = 2$ for all $i \in N$ in the profile $\boldsymbol{\succ}$ of Figure 4. We have that $SW(Plur(\boldsymbol{\succ})) = SW(y) = (n-1)(m-1) = 9$, and that $SW(Plur(\boldsymbol{\succ}^{e,o})) = 1$. This gives us $PLE(Plur, e) \geq 9$, while Lemma 1 only gives us $PLE(Plur, e) \geq 6$.

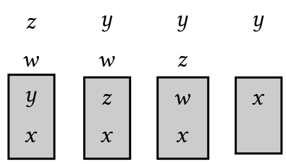

Figure 4: Non-strict PLE lower bound for Plurality.

Our lower bound for Plurality is always exceeded by Borda, under minimal assumptions on the number of voters and alternatives.

Theorem 2. *Let $m \geq 4$ and $n \geq 6$. It holds that $PLE(Plur, e) > PLE(Borda, e)$ for every non-trivial energy function e.*

Proof. Lemma 1 provides a lower bound for Plurality's PLE. It suffices to show that Borda's PLE is always smaller than that. Recall that we define $\alpha = \lfloor \frac{n}{m-1} \rfloor + 1$ if $\lfloor \frac{n}{m-1} \rfloor$ is not an integer, and $\alpha = \frac{n}{m-1}$ otherwise. We claim that $SW(Borda(\boldsymbol{\succ}^{e,o})) > \sum_{i \in \alpha.\,\min_{i\in N} e_i} e_i - 1$ for all $\boldsymbol{\succ}$, e, and o. If this holds, since any possible social welfare is at most $m(n-1)$, the PLE of Borda will smaller than the lower bound for the PLE of Plurality.

To prove our claim, suppose for contradiction that there exist $\boldsymbol{\succ}$, e, and o such that $SW(Borda(\boldsymbol{\succ}^{e,o})) \leq \sum_{i \in \alpha.\,\min_{i\in N} e_i} e_i - 1$ $(*)$. Let $Borda(\boldsymbol{\succ}^{e,o}) = x$. Because x is the Borda winner, every time an alternative $y \neq x$ appears in a ranking within $\boldsymbol{\succ}^{e,o}$ at a position $k \geq 0$ levels above x, it must also appear in some other ranking at a position $\ell \geq k$ levels below x. But if $m \geq 4$ and $n \geq 6$:

$$\alpha \leq \lfloor \frac{n}{m-1} \rfloor + 1 \leq \frac{n}{m-1} + 1 < \frac{n}{m-2} \leq \frac{n}{2}$$

So, from $(*)$ we know that the levels below x in $\boldsymbol{\succ}^{e,o}$ are in total at most $\sum_{\alpha.\,\min_{i\in N} e_i} e_i - 1 < \sum_{\frac{n}{2}.\,\min_{i\in N} e_i} e_i - 1$. This means that the levels above x in $\boldsymbol{\succ}^{e,o}$ are in total strictly more than $\frac{n}{2}(e-1)$, where $e \in \max(\alpha.\,\min_{i\in N} e_i)$. Hence, it is impossible for x to be the Borda winner, and we have reached a contradiction. □

4 THE EFFECTS OF LIMITED ENERGY: AN EXPERIMENTAL ANALYSIS

This section sheds more light to the effects of limited energy on voters. Complementing the analytical methodology of Section 3 that contributed to a qualitative study of the voting rules, we now employ a quantitative approach. Via extensive simulation experiments, we measure the social welfare in multiple voting contexts. In particular, we control a number of crucial parameters:

- The distribution of the intrinsic *preference* profiles.
- The distribution of the *energy limits* of the voters.
- The distribution of the *show orders* applied to the voters.
- The *voting rule* to be used.

To illustrate how the interaction of all the above parameters may affect our results, consider the following example.

Example 2. Suppose we have three alternatives and three voters with identical preferences and energy limits ($e_i = 2$ for all i). What is the optimal distribution of the show orders for the Plurality rule, given that we don't know the voters' exact preferences?

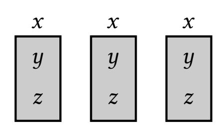

If every voter is shown the alternatives in the same order, then there are three possibilities: voters will all examine the set $\{x, y\}$, or $\{x, z\}$, or $\{y, z\}$ (the last case is depicted above). So, there is 2/3 probability for alternative x (i.e., the alternative offering the highest social welfare) to win; else, y will win. If on the other hand the show order of the alternatives is different for every voter, there is $6/7 > 2/3$ probability for alternative x to be selected: even if one of the voters does not examine it and ranks y first, the other two voters will ensure the optimal plurality winner; this will only fail if one voter is presented with y and then z and another voter is presented with z and then y, causing y to be the plurality winner.

Example 2 targets two extreme cases regarding the discordance between the voters' show orders: no discordance at all (meaning extremely coherent show orders), and as high discordance as possible (meaning totally divergent show orders). The experiments help us examine intermediate values of discordance, as well as various types of preference profiles. Note that when we turn our attention to preference profiles, our focus is similar: we will inspect coherent profiles where voters hold close preferences, but also divergent profiles with contrasting preferences.

4.1 Sampling Method

The sampling method that we employ is based on the *Plackett-Luce* ranking distribution [22, 26, 31] and is "Mallow-like", in the sense of using a reference ranking to which voters' preferences can be more or less similar. This distribution has been found to be practical for simulating choice data [13] and offers high degree of freedom regarding the domains it can model. Our version works as follows.

We first fix a reference ranking $\succ$ over A for the intrinsic preferences of the voters (similarly for the show orders). Let G be a *Gaussian* distribution over the real numbers, with a given mean and entropy. We sample the alternatives appearing in voter i's preference $\succ_i$ sequentially: We draw a value k from G and place on the first level of $\succ_i$ the alternative from the $\lceil k \rceil^{th}$ level of $\succ$ (if $\lceil k \rceil \notin \{1, \ldots, m\}$, we draw again); we then draw another value ℓ and place on the second level of $\succ_i$ the alternative from the $\lceil \ell \rceil^{th}$ level of $\succ$ (if $\lceil \ell \rceil \notin \{1, \ldots, m\}$ or $\lceil \ell \rceil = \lceil k \rceil$, we draw again); we continue until all m alternatives are sampled for $\succ_i$.

More specifically, we consider Gaussian distributions with entropy capturing the degree of discordance between the voters (ranging in $[0, 2]$): high entropy means high discordance.[9] For the voters' energy limits, we draw values from G, with normalised mean in $[0, 1]$. To account for different voting scenarios, we call the profiles (also the show orders and the voters' energy) *coherent*, *similar*, or *divergent* when the entropy of their corresponding G distribution is 0.2, 1, or 1.8, respectively. We call the energy *low*, *medium*, or *high* when the corresponding mean is 0.2, 0.5, and 0.8, respectively.

We sample 3000 preference profiles for each combination of parameters.[10] We discuss results for different numbers of voters, from 10 to 1000, and 5 alternatives. Note that it is not clear whether the number of voters has positive or negative effects on social welfare. For example, take similar profiles and medium coherent energy, with coherent show orders. The social welfare produced by Plurality (similarly by Borda) is on average the same with 10, 100, and 1000 voters, but on large groups we observe more extreme behaviour (Figures 10 and 11 in the Appendix).

In Sections 4.2 and 4.3, our results concern the optimal show orders (♠), the comparative performance of the two voting rules (⋆), and the performance of the voting rules in relation to the discordance between the energy limits of the voters (⧫). These results are easily verifiable in the app; in the interest of space we do not provide relevant figures in the main body of the paper (a screenshot that exemplifies the user interface of the app and the way that our results are demonstrated in it can be seen in Figure 5). Throughout the experiments, we optimise, compare, and contrast the produced (normalised) social welfare values of the two voting rules.

4.2 High/Medium Energy

We discover that in most cases, it is optimal to apply similar show orders to the voters (based on our Gaussian distribution with entropy around 1). See for example Figure 5. This fails in two contexts:[11]

- ♠ For similar or divergent profiles and medium coherent energy, maximising show order entropy is optimal for Borda, especially for larger groups (around 1000 voters).
- ♠ For similar or divergent profiles and medium coherent energy, minimising show order entropy is optimal for Plurality.

To present the comparative performance of the rules, we distinguish two cases. The former concerns all contexts where Plurality is better than Borda, independently of the show orders' distribution.

- ⋆ For fewer voters or divergent profiles, Plurality performs better than Borda. For coherent profiles or for more voters and high energy, Plurality performs close to optimally (see Figure 13 in the Appendix for an illustration) and thus is trivially better than Borda (recall also the insight provided by Proposition 5).
- ⋆ For many voters (around 1000), similar profiles, and medium energy, Plurality always performs better for coherent show orders and Borda always performs better for divergent show orders. For similar show orders, Plurality is better only when we have divergent energy limits.

The next point highlights a basic difference between the two rules.

- ⧫ Plurality performs best for coherent energy, contrary to Borda that for many voters, performs best for divergent energy. See examples in the Appendix, Figures 14 and 15.

4.3 Low Energy

We solve the same exercises for voters with low energy. We see again that similar show orders (linked to entropy around 1) constitute the optimal choice in most cases, except for two contexts for larger groups (around 1000 voters); Figures 16, 17, and 18 in the Appendix exemplify a standard case and the exceptions, respectively:

- ♠ For similar or divergent profiles and coherent energy, maximising the show order entropy is optimal for Borda.
- ♠ For similar or divergent profiles and similar energy, a smaller show order entropy of around 0.5 is optimal for Plurality.[12]

A significant observation is reinforced next: The advantages of Plurality are very prominent for small groups.

- ⋆ For few voters, Plurality's performance is always better than Borda's (or close to it), independently of the show orders.
- ⋆ For many voters (around 1000), Plurality almost always performs better for coherent show orders and Borda for divergent show orders.[13] For similar show orders Plurality is better only when we have divergent energy limits.

Our final remark concerns the comparison between groups with various energy distributions. As opposed to Section 4.2, now it does not hold that coherent energy is beneficial to Plurality.

- ⧫ Both rules perform best for divergent energy, with Plurality exhibiting greater differences than Borda.[14]

4.4 The Big Picture

Here is an overview of our experimental results.

To start, choosing a show order distribution of medium entropy is often the best we can do. The only settings where we should differentiate concern some groups with low to medium energy and non-coherent preferences: the show orders that we should then

[9] The Gaussian distribution with the highest entropy is the uniform distribution.

[10] Conducting recurrent experiments, we observed that sampling 3000 profiles is sufficient to obtain negligible deviations in the results we present.

[11] This can be seen in Figure 12 in the Appendix, which involves the same parameters as Figure 5 except for having coherent instead of divergent energy.

[12] Note the difference with the case of higher energy in Section 4.2, where completely minimising the show order entropy was optimal for Plurality.

[13] The only exceptions being similar or diverging profiles with similar energy against Plurality, and coherent or diverging profiles with diverging energy against Borda.

[14] A minor exception for Plurality is the case of coherent show orders, where groups with divergent energy are not obviously better.

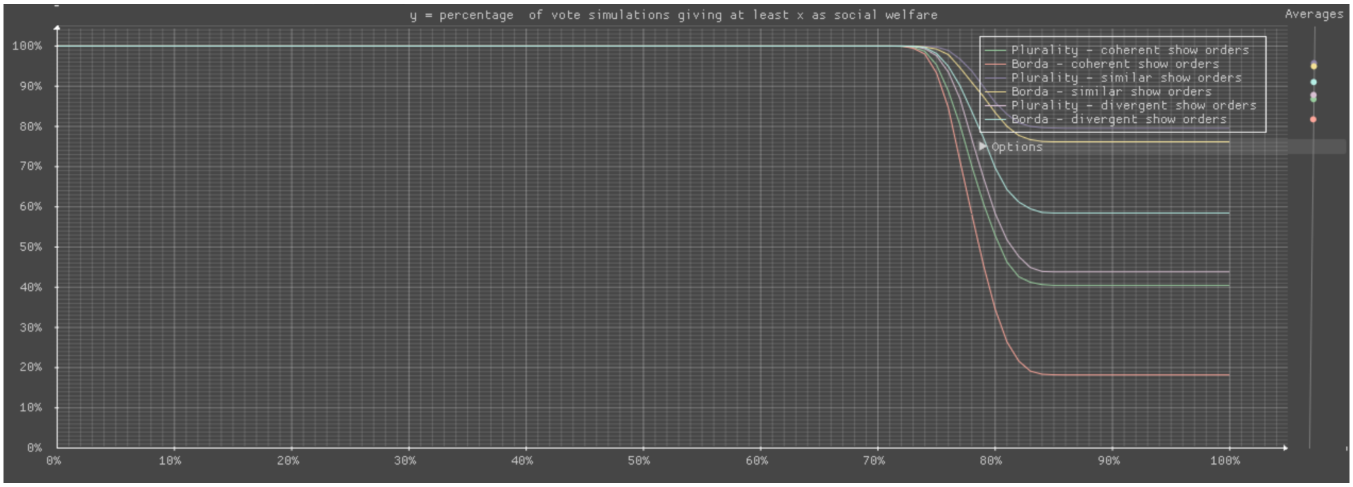

Figure 5: Plurality's and Borda's social welfare for various show orders, similar profiles, medium divergent energy, 1000 voters.

choose are contingent on the voting rule in use. What is also evident is that divergent show orders may only benefit Borda (and are detrimental to Plurality), and *vice versa* for coherent show orders.

Secondly, the performance of Plurality in our framework is especially favourable: For groups with high energy, and even in some cases with medium energy, the collective outcomes that Plurality produces are almost always the ones maximising social welfare. In cases where Plurality is non-optimal, which of the two rules gives better outcomes is an intricate matter of other parameters.

Finally, whether it is profitable to apply each of our rules to groups with coherent or divergent energy depends on the amount of energy that voters have: if their energy is low, groups with divergent energy limits perform better for both rules; otherwise—for high enough energy—groups with coherent energy limits are best for Plurality.

5 STRATEGIC VOTING

Another fundamental aspect of voting methods is their susceptibility to strategic manipulation. So far we have implicitly made a rather simplistic assumption regarding the relationship between the preferences of the voters and their reported votes—namely that they coincide. This is a reasonable starting point for a study in computational social choice; next, it is sensible to move one step forward and wonder whether this is a safe assumption in our framework.

We suppose it is common knowledge that every voter has some energy limit and can compare at least two alternatives, and that the alternatives are presented to every voter in some order. We ask whether voters have an incentive to misrepresent their intrinsic preferences—beyond what is caused by their limited energy—in view of obtaining a favourable outcome. To begin with, let us settle the background. We will build on the standard game-theoretical model of strategic voting in which voters have *full information* about the intrinsic preferences of other members in their group [11, 29], granting additional flexibility with respect to the information they have about the specific energy and show order functions.

To include informational uncertainty, we generalise upon the model of Conitzer et al. [8] and Reijngoud and Endriss [28], also put in use by other works, e.g., [10, 34]. We denote by $\mathcal{I}$ the *information function*: $\mathcal{I}_i(\succ, e, o)$ is the set of profiles of (possibly partial) preferences that voter i deems possible, given a profile of (intrinsic) preferences $\succ$, an energy function e, and a show order function o. An information function $\mathcal{I}$ is *at least as informative* as an information function $\mathcal{I}'$ if $\mathcal{I}_i(\succ, e, o) \subseteq \mathcal{I}'_i(\succ, e, o)$ for all $i, \succ, e, o$.

Note that voter i is only aware of the alternatives $A(e_i, o_i) \subseteq A$. Hence, the profiles she considers possible will only involve this subset. Definition 1 formalises two opposite scenarios.

Definition 1. We say that $\mathcal{I}$ provides:

- *full-energy* and *full-order* information if $\mathcal{I}_i(\succ, e, o) = \{\succ^{e,o}\restriction_{A(e_i,o_i)}\}$, for all $\succ$, e, o, and i
- *zero-energy* and *zero-order* information if $\mathcal{I}_i(\succ, e, o) = \{(\succ^{e',o'}_{-i}\restriction_{A(e_i,o_i)}, \succ_i^{e_i,o_i})$ for all $o' \in \Pi(A)^n$, $e' \in \{2, ..., |A|\}^n\}$, for all $\succ$, e, o, and i

Intuitively, under full-energy and full-order information a voter knows everything about the reported preferences of her peers (conditionally on her own energy and show order); under zero-energy and zero-order information, she knows nothing about them. We define the manipulability of a voting rule in the lines of Reijngoud and Endriss [28], who focus on a *safe* kind of manipulation, i.e., one where the voters report an untruthful preference only if there is no possible profile where this could lead to an inferior outcome.

Definition 2. Consider an energy function e and an information function $\mathcal{I}$. A rule F is *manipulable* on a profile $\succ$ by voter i via $\succ'_i \neq \succ_i^{e,o}$ if for some order function o two conditions hold:

(1) $F(\succ'_{-i}, \succ'_i) \succ_i F(\succ'_{-i}, \succ_i^{e,o})$, for some profile $\succ' \in \mathcal{I}_i(\succ, e, o)$
(2) $F(\succ''_{-i}, \succ_i^{e,o}) \succ_i F(\succ''_{-i}, \succ'_i)$, for no profile $\succ'' \in \mathcal{I}_i(\succ, e, o)$

Since voters are allowed to submit partial preferences, manipulation may be of two different kinds: voters may either *omit* pairwise comparisons between alternatives they have examined, or *flip* their truthful preferences. We borrow the following definitions [20]:[15]

Definition 3. We say that $\succ'$ is induced from $\succ$ by:

- *omission* when for all $x, y \in A$, $x \succ' y$ implies that $x \succ y$;
- *flipping* when for all $x, y \in A$, $x \succ' y$ holds if and only if $x \succ y$ or $y \succ x$.

We say that a voting rule is *manipulable* by omission (flipping) if it is manipulable on some profile $\boldsymbol{\succ}$ by voter i via $\succ_i' \neq \succ_i^{e,o}$ and $\succ_i'$ is induced from $\succ_i^{e,o}$ by omission (flipping).

Lemma 2 verifies a natural relation of information functions—it is a direct analogue of Lemma 1 by Reijngoud and Endriss [28].

Lemma 2. *If $\mathcal{I}$ is at least as informative as $\mathcal{I}'$, then a voting rule that is manipulable under $\mathcal{I}'$ will also be manipulable under $\mathcal{I}$.*

The next two theorems characterise the energy limits instigating manipulability. Plurality is susceptible to manipulation even under very low energy—this is not the case for Borda, which requires a slightly higher amount of energy to be manipulated by omission.

Theorem 3. *Under any information function, the Plurality rule is manipulable for an energy function e by omission (similarly by flipping) if and only if $e_i \geq 3$ for some voter i.*

Proof. First, with $e_i < 3$ for all i, no voter has an incentive to manipulate. So consider an arbitrary energy function e and suppose without loss of generality that $e_1 \geq 3$. Take $\mathcal{I}$ to be the zero-energy and zero-order information function. We will show that Plurality is manipulable for $\mathcal{I}$ and the statement will follow from Lemma 2.

Take a two-voter profile $\boldsymbol{\succ}$ and a function o such that $\boldsymbol{\succ}\restriction_{A(e_1,o_1)} = (\succ_1, \succ_2)$, where $x_3 \succ_1 x_1 \succ_1 x_2$ and $x_2 \succ_2 x_1 \succ_2 x_3$. We have that $\boldsymbol{\succ}\restriction_{A(e_1,o_1)} \in \mathcal{I}_1(\boldsymbol{\succ}, e, o)$ and $Plur(\boldsymbol{\succ}\restriction_{A(e_1,o_1)}) = x_2$. Under zero-energy and zero-order information, voter 1 is aware of the intrinsic preference $\succ_2$, but only knows that $e_2 \geq 2$. This guarantees her that x_3 cannot be the Plurality winner in $\boldsymbol{\succ}^{e,o}$. By submitting the untruthful partial preference $x_1 \succ_1' x_2$ instead of her truthful one, voter 1 has nothing to lose—yet, she can ensure a preferable outcome under some possible profile because $Plur(\succ_1', \succ_2) = x_1$. Similarly, voter 1 could flip her preference to $x_1 \succ_1' x_2 \succ_1' x_3$. □

Theorem 4. *Under any information function, the Borda rule is manipulable for an energy function e by flipping if and only if $e_i \geq 3$ for some voter i.*

Proof. Consider $\boldsymbol{\succ}\restriction_{A(e_1,o_1)} = (\succ_1, \succ_2)$, where $x_2 \succ_1 x_1 \succ_1 x_3$ and $x_1 \succ_2 x_2 \succ_2 x_3$. As in the proof of Theorem 3, under any energy and show order information voter 1 knows that x_3 cannot be the Borda winner: she can manipulate with $x_2 \succ_1' x_3 \succ_1' x_1$. □

Theorem 5. *Under any information function, the Borda rule is manipulable for an energy function e by omission if and only if $e_i \geq 5$ for some voter i.*

Proof. Suppose first that $e_i \leq 4$ for all voters i, and consider $\succ_i'$ to be induced from $\succ_i$ by omission. From the contrapositive of Lemma 2, it suffices to show that the Borda rule is not manipulable under full-energy and full-order information. Precisely because $|A(e_i, o_i)| \leq 4$ for all $i \in N$, and by definition of the Borda rule for partial preferences, we have that $s^{Borda}(\succ_i)(x) - s^{Borda}(\succ_i)(y) \geq s^{Borda}(\succ_i')(x) - s^{Borda}(\succ_i')(y)$ for all alternatives x, y such that $x \succ_i y$. Theorem 2 of Kruger and Terzopoulou [20] then tells us that voter i cannot manipulate by reporting $\succ_i'$ instead of $\succ_i$.[16]

Then, suppose that $e_i \geq 5$ for some i and take $\mathcal{I}$ to be the zero-energy and zero-order information function. We will show that Borda is manipulable for $\mathcal{I}$, so by Lemma 2 it is manipulable for every information function. Consider a two-voter profile $\boldsymbol{\succ}$ and a function o such that $\boldsymbol{\succ}\restriction_{A(e_1,o_1)} = (\succ_1, \succ_2)$, where $x_2 \succ_1 x_1 \succ_1 x_5 \succ_1 x_4 \succ_1 x_3$ and $x_1 \succ_2 x_2 \succ_2 x_3 \succ_2 x_4 \succ_2 x_5$. We have that $\boldsymbol{\succ}\restriction_{A(e_1,o_1)} \in \mathcal{I}_1(\boldsymbol{\succ}, e, o)$ and $Borda(\boldsymbol{\succ}\restriction_{A(e_1,o_1)}) = x_1$. Independently of voter 2's energy and the order in which the alternatives are presented to her, it is never harmful for voter 1 to submit the untruthful partial preference $\succ_1' = x_2 \succ_1 x_5 \succ_1 x_4 \succ_1 x_3$ instead of her truthful one. It may even be beneficial: $Borda(\succ_1', \succ_2) = x_2$. □

Note that the informational assumptions of this section are consequences of limited energy: although the voters know the intrinsic preferences of their peers to some extent, they are restricted both with respect to the alternatives of which they are aware themselves and the votes of the other members in their group that result from the applied show orders. To the best of our knowledge, this set of assumptions has not been considered before in strategic voting, and our results cannot be reproduced in existing frameworks. Nonetheless, we do enforce a large pool of literature in social choice, indicating that strategic behavior is often unavoidable.

6 CONCLUSION

We present a framework for limited energy in voting. We formalise two important notions—the voters' energy limits and the order in which the alternatives are presented to them—and focus on the two most prominent scoring rules: Plurality and Borda. We provide a foundational analysis of social welfare, conduct extensive simulation experiments, and also explore the strategic incentives of voters with limited energy.

We stress a number of take-home messages: With small groups, high energy, or coherent show orders, Plurality should be preferred to Borda; but Borda is better with large groups and divergent show orders. Under low energy, groups with divergent energy perform better. Overall, the fact that Plurality is less sensitive to loss of information is advantageous in many contexts. The Borda rule however is slightly more immune to strategic behaviour, demanding more energy from those that manipulate by omission.

This paper paves the way for intriguing research outside its current scope. Indicatively: How robust are our results for different notions of social welfare or different voting rules? What is the computational complexity of discovering the optimal show order given a profile distribution? Is it easy to compute the PLE value for different energy functions, and if not, are there desirable approximations? Which are the exact conditions on the energy and the preferences of the voters under which a mechanism designer can control the voting outcome by selecting suitable show orders? Possible directions are plentiful.

[15] Note that Kruger and Terzopoulou [20] also consider manipulation by *adding* pairwise comparisons in a voter's preference. In our model this is not meaningful, since the voters are not aware of any alternatives beyond the ones they have examined.

[16] This is rather intuitive: Not being able to increase the score difference between a more preferred and a less preferred alternative, a voter has no incentive to manipulate.

7 ACKNOWLEDGMENTS

I am deeply grateful to the SpirOps-AI company in Paris, in particular to Edouard Emberger and Romane Flechard, for their tremendous help in the development of the app for voting with limited energy. This collaboration would not be possible without the Paris Region Fellowship Programme (grant agreement No. 945298). Thanks also to Simon Rey and Arianna Novaro for their valuable feedback during the writing of this paper, to the anonymous AAMAS reviewers, and to the audience of the Theory seminar at the Athens University of Economics and Business where a preliminary version of this work was presented.

REFERENCES

[1] Manel Ayadi, Nahla Ben Amor, and Jérôme Lang. 2022. Approximating Voting Rules from Truncated Ballots. *Autonomous Agents and MultiAgent Systems* 36, 1 (2022), 1–29.

[2] Nawal Benabbou, Serena Di Sabatino Di Diodoro, Patrice Perny, and Paolo Viappiani. 2016. Incremental Preference Elicitation in Multi-Attribute Domains for Choice and Ranking with the Borda Count. In *Proceedings of the 10th International Conference on Scalable Uncertainty Management (SUM)*.

[3] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

[4] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D Procaccia. 2013. How Bad is Selfish Voting?. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*.

[5] Guillaume Cabanac and Thomas Preuss. 2013. Capitalizing on Order Effects in the Bids of Peer-Reviewed Conferences to Secure Reviews by Expert Referees. *Journal of the American Society for Information Science and Technology* 64, 2 (2013), 405–415.

[6] Ioannis Caragiannis, George Krimpas, and Alexandros Voudouris. 2015. Aggregating Partial Rankings with Applications to Peer Grading in Massive Online Open Courses. In *Proceedings of the 14th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[7] Vincent Conitzer and Tuomas Sandholm. 2005. Communication Complexity of Common Voting Rules. In *Proceedings of the 6th ACM conference on Electronic Commerce (EC)*.

[8] Vincent Conitzer, Toby Walsh, and Lirong Xia. 2011. Dominating Manipulations in Voting with Partial Information. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*.

[9] Carmel Domshlak, Eyke Hüllermeier, Souhila Kaci, and Henri Prade. 2011. Preferences in AI: An Overview. *Artificial Intelligence* 175, 7–8 (2011), 1037–1052.

[10] Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein. 2016. Strategic Voting with Incomplete Information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*.

[11] Allan Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 4 (1973), 587–601.

[12] Gerd Gigerenzer and Reinhard Selten. 2002. *Bounded Rationality: The Adaptive Toolbox*. MIT press.

[13] Isobel Claire Gormley and Thomas Brendan Murphy. 2005. *Exploring Heterogeneity in Irish Voting Data: A Mixture Modelling Approach*. Technical Report. Technical Report, Department of Statistics, Trinity College Dublin.

[14] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. 2013. Restricted Manipulation in Iterative Voting: Condorcet Efficiency and Borda Score. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*.

[15] Aleksandar Hatzivelkos. 2018. Borda and Plurality Comparison with Regard to Compromise as a Sorites Paradox. *Interdisciplinary Description of Complex Systems: INDECS* 16, 3B (2018), 465–484.

[16] Meir Kalech, Sarit Kraus, Gal A Kaminka, and Claudia V Goldman. 2011. Practical Voting Rules with Partial Information. *Autonomous Agents and Multi-Agent Systems* 22, 1 (2011), 151–182.

[17] Kathrin Konczak and Jérôme Lang. 2005. Voting Procedures with Incomplete Preferences. In *Proceedings of the IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling*.

[18] Aaron Koolyk, Tyrone Strangway, Omer Lev, and Jeffrey S Rosenschein. 2017. Convergence and Quality of Iterative Voting under Non-Scoring Rules. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*.

[19] Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-Case Equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*. 404–413.

[20] Justin Kruger and Zoi Terzopoulou. 2020. Strategic Manipulation with Incomplete Preferences: Possibilities and Impossibilities for Positional Scoring Rules. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems(AAMAS)*.

[21] Tyler Lu and Craig Boutilier. 2011. Robust Approximation and Incremental Elicitation in Voting Protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*.

[22] R Duncan Luce. 1959. *Individual Choice Behavior: A Theoretical Analysis*. Wiley.

[23] Reshef Meir. 2017. Iterative Voting. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access.

[24] Reshef Meir, Omer Lev, and Jeffrey S Rosenschein. 2014. A Local-Dominance Theory of Voting Equilibria. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*.

[25] Maria Silvia Pini, Francesca Rossi, K. Brent Venable, and Toby Walsh. 2008. Aggregating Partially Ordered Preferences. *Journal of Logic and Computation* 19, 3 (2008), 475–502.

[26] Robin L Plackett. 1975. The Analysis of Permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 24, 2 (1975), 193–202.

[27] Ariel D Procaccia and Jeffrey S Rosenschein. 2006. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*.

[28] Annemieke Reijngoud and Ulle Endriss. 2012. Voter Response to Iterated Poll Information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[29] Mark Allen Satterthwaite. 1975. Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory* 10, 2 (1975), 187–217.

[30] Esther-Mirjam Sent. 2018. Rationality and Bounded Rationality: You Can't Have One Without the Other. *The European Journal of the History of Economic Thought* 25, 6 (2018), 1370–1386.

[31] Arthur Richard Silverberg. 1980. *Statistical Models for q-Permutations*. Ph.D. Dissertation. Princeton University.

[32] Herbert A Simon. 1956. Rational Choice and the Structure of the Environment. *Psychological review* 63, 2 (1956), 129–138.

[33] Zoi Terzopoulou. 2021. *Collective Decisions with Incomplete Individual Opinions*. Ph.D. Dissertation. Institute for Logic, Language and Computation, University of Amsterdam.

[34] Yuliya A Veselova. 2020. Does Incomplete Information Reduce Manipulability? *Group Decision and Negotiation* 29, 3 (2020), 523–548.

[35] William S Zwicker. 2016. Introduction to the Theory of Voting. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press.