

On the Welfare of Cardinal Voting Mechanisms

Umang Bhaskar¹

Tata Institute of Fundamental Research, Mumbai, India
umang@tifr.res.in

Abheek Ghosh

The University of Texas at Austin, TX, USA
abheek@cs.utexas.edu

Abstract

A voting mechanism is a method for preference aggregation that takes as input preferences over alternatives from voters, and selects an alternative, or a distribution over alternatives. While preferences of voters are generally assumed to be cardinal utility functions that map each alternative to a real value, mechanisms typically studied assume coarser inputs, such as rankings of the alternatives (called ordinal mechanisms). We study cardinal mechanisms, that take as input the cardinal utilities of the voters, with the objective of minimizing the distortion – the worst-case ratio of the best social welfare to that obtained by the mechanism.

For truthful cardinal mechanisms with m alternatives and n voters, we show bounds of $\Theta(mn)$, $\Omega(m)$, and $\Omega(\sqrt{m})$ for deterministic, unanimous, and randomized mechanisms respectively. This shows, somewhat surprisingly, that even mechanisms that allow cardinal inputs have large distortion. There exist ordinal (and hence, cardinal) mechanisms with distortion $O(\sqrt{m \log m})$, and hence our lower bound for randomized mechanisms is nearly tight. In an effort to close this gap, we give a class of truthful cardinal mechanisms that we call randomized hyperspherical mechanisms that have $O(\sqrt{m \log m})$ distortion. These are interesting because they violate two properties – localization and non-perversity – that characterize truthful ordinal mechanisms, demonstrating non-trivial mechanisms that differ significantly from ordinal mechanisms.

Given the strong lower bounds for truthful mechanisms, we then consider approximately truthful mechanisms. We give a mechanism that is δ -truthful given $\delta \in (0, 1)$, and has distortion close to 1. Finally, we consider the simple mechanism that selects the alternative that maximizes social welfare. This mechanism is not truthful, and we study the distortion at equilibria for the voters (equivalent to the Price of Anarchy, or PoA). While in general, the PoA is unbounded, we show that for equilibria obtained from natural dynamics, the PoA is close to 1. Thus relaxing the notion of truthfulness in both cases allows us to obtain near-optimal distortion.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory

Keywords and phrases computational social choice, voting rules, cardinal mechanisms, price of anarchy, distortion

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2018.27

1 Introduction

A society or a group of people may have different views and preferences but want to make a collective decision that will impact the entire group. For example, the people of India may have conflicting opinions on which party should win the Lok Sabha elections, and who should be the Prime Minister. This is the problem of preference aggregation, and the methods

¹ Research funded in part by a Ramanujan fellowship.



© Umang Bhaskar and Abheek Ghosh;

licensed under Creative Commons License CC-BY

38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2018).

Editors: Sumit Ganguly and Paritosh Pandya; Article No. 27; pp. 27:1–27:22

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of achieving this aggregation are called voting mechanisms – functions that map the given preferences of voters over a set of alternatives to a single alternative or a distribution over alternatives, without money being exchanged.

Central to the question of preference aggregation is the question of how preferences are perceived by the voters, and how they are expressed to the mechanism. In classical social choice theory, particularly when a voting mechanism is randomized, i.e., can output a distribution over the set of alternatives, voter preferences are assumed to be von Neumann-Morgenstern utility functions, that map each alternative to a real-valued utility. We assume that the total utility each voter has for the alternatives is 1. This is the unit-sum assumption, though other normalizations such as unit-range are also studied. A voter then prefers distributions over the alternatives that maximize her expected utility. These utility functions may be latent and hidden from the mechanism but are required for the voter to rationally compare distributions or lotteries over the set of outcomes. In contrast to these cardinal utility functions of the voters, frequently the mechanisms studied in the literature, and used in practice, have coarser inputs, such as a ranking of the alternatives, or simply a vote for the alternative with the highest utility (called ordinal and plurality voting respectively).

If utility functions are cardinal, then an understanding of cardinal voting mechanisms, where voters give as input their utility functions, seems fundamental to understanding the problem of preference aggregation. Though less popular than other voting mechanisms owing to their complexity, cardinal voting mechanisms find use in many areas. For example, they are motivated by automated agents in recommender systems that use exact numeric values for making decisions, and hence naturally have easily expressible cardinal utilities. The use of these automated agents in a movie recommendation system is described by Ghosh et al. [14] (cf. [24]). Hillinger further argues for the use of cardinal voting mechanisms, especially since they do not artificially restrict the freedom of expression of voters [17].

Given an input format for voter preferences, how then should the mechanism choose an alternative? A widely studied property is incentive compatibility or truthfulness – a voter should maximize her expected utility by truthfully expressing her preferences, irrespective of the votes of others. Truthful mechanisms are desirable since voters need not strategize or seek information on the behavior of other voters. Other properties for mechanisms that are studied include Pareto-efficiency and polynomial-time computation. A natural objective for the voting mechanism, given the cardinal utilities of voters, is to maximize the social welfare or the total utility of the voters. This has been a mechanism objective in a number of recent papers (e.g., [7, 13]). The objective of social welfare assumes that the utilities allow for interpersonal comparison: that a unit of voter 1’s utility is equivalent to a unit of agent 2’s utility. Such comparisons may not be generally applicable, but even then, aggregate utility (or disutility) is frequently used as a quantitative measure, e.g., man-hours required for a project, or total time spent in traffic. The motivation for studying social welfare from classical economic theory, as well as further uses in modern recommendation systems, is also described by Boutilier et al. [7].

The social welfare of a voting mechanism is measured by the distortion – the ratio of the welfare of the best alternative to the expected welfare obtained by the mechanism, in the worst case over all instances [24].² Unfortunately, combined with the requirement that the

² In their paper, Procaccia and Rosenschein use distortion to measure the loss due to the embedding of cardinal utilities into the space ordinal preferences. However, they also define distortion as stated here. Other papers use ‘approximation ratio’ for this quantity (e.g., [13]). We find the term distortion to be more natural and descriptive, and hence use it here for cardinal mechanisms as well.

mechanism be truthful, for the input formats typically studied, the distortion is known to be large – for ordinal mechanisms, that take as input a ranking of the alternatives, the distortion is $\Theta(\sqrt{m \log m})$ [5], where m is the number of alternatives. Even for ordinal mechanisms that are not truthful, the distortion is $\Omega(\sqrt{m})$ [7], suggesting that most of the loss in welfare is caused by the incongruity between the cardinal utilities experienced by the voters, and ordinal preferences given as input to the mechanism. This begs the question: for truthful, cardinal mechanisms, where the input to the mechanisms are cardinal utility functions, how large is the distortion? Note that in this case, there is no incongruity between the preferences experienced and expressed: both are cardinal.

Our goal in this paper is to address this question, and to obtain bounds on the social welfare obtainable by cardinal mechanisms. While cardinal mechanisms are less popular than mechanisms with simpler input formats, we believe that an understanding of cardinal mechanisms is crucial in many ways to understanding other preference aggregation mechanisms. Firstly, lower bounds obtained for cardinal mechanisms are lower bounds for mechanisms with other input formats as well. Secondly, studying cardinal mechanisms helps disentangle the effects of various constraints on the mechanism, since the input format is no longer a constraint. Thirdly, as noted above, particularly in the case of automated agents, cardinal mechanisms are of practical use. Lastly, we note that cardinal mechanisms have received less attention than ordinal mechanisms, and theoretically these present several challenging questions. For example, strong characterizations of truthful ordinal mechanisms are known [16, 1], while truthful cardinal mechanisms are only partially characterized (e.g., [18, 3]). This is a long-standing open question, and we hope that the perspective of distortion of cardinal mechanisms may present useful insights for this question as well.

We study truthful, nearly-truthful, as well as manipulable cardinal mechanisms, and provide bounds on the social welfare obtainable. In the last case, we study the social welfare at equilibrium for the deterministic cardinal mechanism that simply returns the alternative with maximum social welfare for the reported utility functions. While for truthful cardinal mechanisms we show strong and nearly tight lower bounds, for the last two cases we are able to show that with some reasonable restrictions, the distortion achieved is nearly 1, in sharp contrast to the case of truthful cardinal mechanisms, where the distortion is $\Omega(\sqrt{m})$.

Our Contribution. In this paper, we study the distortion of cardinal mechanisms with n unit-sum voters and $m \geq 3$ alternatives. Unit-sum assumes that the sum of utilities of each voter for the alternatives is 1. This assumption implies that all the voters have equal weight, and no voter is more important than the other in contributing to the social welfare. We first focus on truthful cardinal mechanisms, and show that for deterministic cardinal mechanisms, the distortion is $\Theta(mn)$. Note that the trivial randomized mechanism that picks an alternative uniformly without looking at the utilities of the voters has distortion $O(m)$. A natural property for mechanisms is Pareto-optimality (e.g., [16]), which states that for any alternative a chosen by the mechanism with positive probability, there is no alternative b for which all voters have higher utility. A significantly weaker property is unanimity, which states that if there is an alternative that has maximum utility for all voters, then the mechanism should pick this with probability 1. Unfortunately, even for this weaker property, we show that any truthful unanimous mechanism has distortion $\Omega(m)$. Underlying our results are strong previous characterizations of unanimous mechanisms which show that such mechanisms must be random dictatorships, which pick a voter at random and return the maximum utility alternative for the voter [10, 18, 21].

We then focus on truthful randomized mechanisms. There exist randomized ordinal (and hence cardinal) mechanisms with distortion $O(\sqrt{m \log m})$. We show that, perhaps surprisingly, this is nearly the best possible for cardinal mechanisms as well, showing a bound of $\Omega(\sqrt{m})$ on the distortion of cardinal mechanisms. Note that if voters reported their utility truthfully, a mechanism with distortion 1 is trivial. This implies that the loss in welfare due to truthfulness is already large, apart from that caused by the information loss due to the input format.³

We leave the problem of closing the gap between the upper and lower bounds ($O(\sqrt{m \log m})$ and $\Omega(\sqrt{m})$) open. We instead address the following question: can cardinal mechanisms with low distortion be very different from ordinal mechanisms? In particular, the characterization by Gibbard of truthful ordinal mechanisms requires such mechanisms to satisfy non-perversity and localization (these are defined in the next section) [16]. Must mechanisms that violate these properties be trivial, with large distortion, or do there exist cardinal mechanisms that violate these properties, and yet have good distortion? We give a mechanism that we call a randomized hyperspherical mechanism that violates these properties, but has distortion $O(\sqrt{m \log m})$, matching the best known upper bound. Spherical mechanisms were previously studied by Feige and Tennenholtz [12], but these are much simpler mechanisms with distortion $\Omega(m)$. We view our mechanism as a significant extension of these. The mechanism we introduce may be of independent technical interest as well. One of the steps involved in the mechanism is to project a point onto the intersection of the standard simplex and a hypersphere of given radius, for which we give a polynomial time algorithm.

Given the strong lower bounds, we then consider two kinds of mechanisms that may incentivize strategic behavior (called manipulable). We first study approximately-truthful cardinal mechanisms, where a mechanism is δ -truthful if no voter can increase her expected utility by more than δ by reporting her utilities untruthfully. Here we show surprisingly good results: for any $\delta \in (0, 1)$, we give a cardinal mechanism that has distortion that approaches 1 as the number of voters increases, and is 2δ -truthful. Thus slightly relaxing the notion of truthfulness allows us to obtain near optimal bounds on the distortion. It is instructive to compare our results with those of Birrell and Pass, who show that approximately truthful ordinal mechanisms can be used to approximate any deterministic ordinal mechanism, in a formal sense [6]. However, the distortion of any deterministic ordinal mechanism is $\Omega(m)$ [24].

We lastly consider the natural, but manipulable, deterministic mechanism that for any utilities given as input, simply returns the alternative that has maximum social welfare according to these utilities. Since we can no longer rely on voters being truthful, we instead consider the pure Nash equilibrium for this mechanism, i.e., utility profiles where no single voter can report a different utility function and improve her utility. Here, the distortion is equivalent to the Price of Anarchy (PoA), defined as the ratio of the welfare of the best alternative to the expected welfare of the mechanism, in the worst case over all equilibria and all instances. Simple examples show the PoA is in general unbounded, and even natural refinements studied previously have unbounded PoA. Instead, we consider equilibria reachable from natural iterative voting, where in each step, a voter changes her input to the mechanism to improve her utility. Iterative voting has been considered in a number of previous papers (e.g., [20, 25]), though usually for ordinal or other mechanisms with restricted input formats. We show that under certain natural restrictions on the allowable deviations, iterative voting converges, and the PoA approaches 1 as the number of voters increases.

³ Ariel Procaccia in a conversation mentioned that he had also obtained this bound independently, but had not written it up.

Related Work. We focus here on literature directly related to our work, and refer to the book on computational social choice [8] and the survey by Barbera [2] for a wider discussion. For voting mechanisms, truthfulness has been an important consideration, epitomized by the well-known impossibility result by Gibbard and Satterthwaite for deterministic ordinal voting mechanisms which states that any such truthful mechanism must be a dictatorship [15, 27]. Allowing randomization alleviates this, and such mechanisms were characterized by Gibbard [16] as being distributions over unilateral and dupe mechanisms, each of which must be localized and non-perverse. This characterization was further refined by Barbera [1].

In comparison, for characterizing truthful cardinal mechanisms, only partial results are known. Characterizations are known for twice-differentiable truthful cardinal mechanisms [3]. Further, truthful cardinal mechanisms which are unanimous are a convex combination of dictatorial schemes [18, 10, 21].

The social welfare of voting mechanisms was explicitly studied by Boutilier et al. [7], who showed that for randomized ordinal mechanisms, the distortion was $O(\sqrt{m} \log^* m)$ and $\Theta(\sqrt{m})$. The concept of distortion as a measure of loss of welfare by a mechanism was introduced earlier and studied for many well-known deterministic voting mechanisms including Borda, Plurality, and Veto [24]. In general, the distortion is shown to be unbounded. For highly structured utility functions, the authors obtain positive results. The mechanisms studied in these papers are not truthful. For truthful ordinal mechanisms for voters with unit-sum utilities, the distortion is known to be $\Theta(\sqrt{m} \log m)$ [5].

Filos-Ratsikas and Miltersen study the social welfare of truthful cardinal mechanisms with voters that have unit-range utilities, i.e., for each voter i , $\max_{a \in A} u_i(a) = 1$ and $\min_{a \in A} u_i(a) = 0$ [13]. They obtain bounds of $O(m^{3/4})$ and $\Omega(\log \log m / \log m)$ for truthful cardinal mechanisms, and a bound of $\Omega(m^{2/3})$ for ordinal mechanisms and some generalizations.

Relaxing truthfulness, it is known that approximately truthful randomized ordinal mechanisms can approximate any deterministic ordinal voting rule, in that the output obtained by the randomized mechanism could have been obtained by the deterministic ordinal voting rule by changing a small number of votes [6]. This does not give us bounds on the distortion, since any deterministic ordinal mechanism has distortion $\Omega(m)$ (e.g., [24]).

The social welfare at equilibrium – with the ratio well-known as the Price of Anarchy (PoA) – is known to be bad if the set of equilibria is not restricted. Hence a number of papers study equilibria reachable by natural best-response dynamics, called iterative voting. For the plurality voting rule, best-response dynamics is shown to converge in $O(mn)$ steps [20]. Convergence for Veto, Borda, and other voting rules is also studied [19, 26]. Further, the social welfare for equilibria obtained through best-response dynamics for plurality is known to be at most 1 less than the optimal, however, this can be small (and hence the PoA can be large) for Veto and Borda. Finally, further restrictions on the class of equilibria obtainable have also been studied for plurality voting, such as strong equilibria or equilibria with truth-biased voters [11, 23]. In particular, Rabinovich et al. [25] characterize the set of equilibria obtainable from plurality voting with truth-biased voters.

2 Preliminaries and Notation

A population of n voters (or agents) $N = \{1, 2, \dots, n\}$ want to select an alternative (or candidate) from a set A of size m . Each voter i has a utility function $u_i : A \rightarrow [0, 1]$. We will sometimes abuse notation and think of u_i as vector in \mathbb{R}^m . We assume that the voters

are unit-sum voters, i.e., the total utility of a voter is 1 ($\sum_{a \in A} u_i(a) = 1$ for each voter i). We will also assume that no two alternatives have the same utility for a voter, thus if $u_i(a) = u_i(b)$, then $a = b$. This is a mild technical assumption which allows us to obtain a total order over alternatives for each voter, and simplify our proofs. Let $\vec{u} := (u_i)_{i \in N}$ be a vector of utility functions, or a utility profile. Define $\vec{u}/_i u'_i$ to be the utility profile \vec{u} with the utility function for the i th voter replaced by u'_i .

A (cardinal) mechanism μ is defined as a map, possibly randomized, from input utility profiles to distributions over alternatives. For our lower bounds, we assume that the mechanisms we study have at least 3 feasible alternatives, i.e., there exist three alternatives each of which is chosen with positive probability for some input utility profile. We will sometimes compare our mechanisms to ordinal mechanisms, which are similarly possibly randomized maps from rankings over the set of alternatives provided by voters, to distributions over the alternatives. A plurality mechanism is one where each voter selects a single alternative, and the mechanism chooses an alternative that is selected by the maximum number of voters. Given a distribution $(p_a)_{a \in A}$ over the alternatives, the expected utility of a voter is $\sum_{a \in A} u_i(a) p_a$. We assume that all voters are expected utility maximizers, and note that to maximize her utility a voter i may report a different utility function $u'_i \neq u_i$ to the mechanism. For clarity, we call the input provided to a mechanism by a voter her strategy, which may not be her true utility. A mechanism is truthful (more formally, truthful in expectation) if each voter obtains maximum expected utility by reporting her true utility profile u_i , irrespective of the strategies of the other voters. Thus a cardinal mechanism is truthful if for each voter i with utility u_i , and each strategy profile \vec{u}' , $\mathbb{E}_{a \sim \mu(\vec{u}')} [u_i(a)] \leq \mathbb{E}_{a \sim \mu(\vec{u}'/_i u_i)} [u_i(a)]$. A mechanism is manipulable if it is not truthful. Further, relaxing truthfulness, in a δ -truthful mechanism each voter can improve her expected utility by at most δ by choosing a strategy that is not her true utility. In both these cases, we assume that voters vote truthfully, since the incentive to misreport utility functions is small.

For an alternative a in an instance with utility profile \vec{u} , the utilitarian social welfare (or simply welfare) is $\text{sw}(a) = \sum_i u_i(a)$, the sum of utilities of all the voters for that alternative. We study mechanisms that maximize welfare, and hence our primary measure of a mechanism is its distortion, defined informally as the worst-case ratio of the maximum utility of an alternative, to the expected utility obtained by the mechanism [4, 24]. The distortion of a mechanism is thus at least 1. Under the assumption that voters choose their utility u_i as strategy, the distortion for a mechanism μ is defined formally as:

$$\text{dist}(\mu) := \sup_{\vec{u}} \frac{\max_{a \in A} \text{sw}(a)}{\mathbb{E}_{a \sim \mu(\vec{u})} \text{sw}(a)}.$$

In Section 5, we study the simple deterministic mechanism that chooses the alternative with maximum welfare, for the strategy profile reported by the voters. This mechanism is not truthful, and we will be interested in welfare at the pure Nash equilibrium for the voters. In this case, the distortion is equivalent to the well-studied Price of Anarchy. Formally, let \vec{u} be the utility profile of the voters, and let \vec{u}' be a strategy profile. Then \vec{u}' is a pure Nash equilibrium if for every voter $i \in N$ and every strategy u''_i ,

$$\mathbb{E}_{a \sim \mu(\vec{u}')} u_i(a) \geq \mathbb{E}_{a \sim \mu(\vec{u}'/_i u''_i)} u_i(a).$$

The Price of Anarchy (PoA) of a mechanism μ is defined as the worst-case ratio over all possible instances, of the maximum welfare of an alternative to the lowest welfare of an alternative chosen by the mechanism at an equilibrium.

$$\sup_{\vec{u}} \sup_{\text{equilibrium strategy profiles } \vec{u}'} \frac{\max_{a \in A} \sum_{i \in N} u_i(a)}{\mathbb{E}_{a \sim \mu(\vec{u}')} \sum_{i \in N} u_i(a)}$$

The following properties for ordinal mechanisms were introduced by Gibbard [16].

► **Definition 1.** An ordinal mechanism is non-perverse if increasing the ranking of an alternative by a voter, while leaving the relative order of the other alternatives unchanged, does not decrease the probability that the alternative is selected. An ordinal mechanism is localized if, for two rankings by a voter where the set of top k alternatives are the same (but individual rankings may not be preserved), the total probability mass on these k alternatives is also unchanged.

We extend these properties to cardinal mechanisms as follows.

► **Definition 2.** Let utility functions v, v' be such that $v'(a) > v(a)$ for some alternative a , and for all other alternatives $b, c \neq a$, $v(b) \geq v(c)$ iff $v'(b) \geq v'(c)$. That is, the relative order of the other alternatives remains unchanged. Cardinal mechanism μ is non-perverse if for all utility profiles \vec{u} and voters i , $\mu(\vec{u}/_i v')(a) \geq \mu(\vec{u}/_i v)(a)$.

► **Definition 3.** Given utility functions v, v' and a permutation π so that $v(a_1) \geq \dots \geq v(a_m)$ and $v'(a_{\pi(1)}) \geq \dots \geq v'(a_{\pi(m)})$, cardinal mechanism μ is localized if for every $k \leq m$ such that (i) $\{a_1, \dots, a_k\} = \{a_{\pi(1)}, \dots, a_{\pi(k)}\}$ and (ii) $\sum_{i=1}^k v(a_i) = \sum_{i=1}^k v(a_{\pi(i)})$, the probability mass on these k alternatives remains unchanged. That is, for every utility profile \vec{u} and voter i , $\sum_{i=1}^k \mu(\vec{u}/_i v)(a_i) = \sum_{i=1}^k \mu(\vec{u}/_i v')(a_{\pi(i)})$.

All missing proofs are given in the appendix.

3 Truthful Mechanisms

In this section, we obtain bounds on the distortion of truthful mechanisms. Disappointingly, but perhaps unsurprisingly, we first show that deterministic mechanisms have distortion $\Theta(mn)$ (Theorem 7). Further, mechanisms that are unanimous – i.e., if there exists an alternative that has maximum utility for each voter, then this alternative must be selected – have distortion $\Omega(m)$ (Theorem 9).⁴ In fact, we show that truthfulness is in general expensive – all truthful cardinal mechanisms, even randomized, have distortion $\Omega(\sqrt{m})$ (Theorem 10). The lower bound is disappointing, since it shows that even truthful mechanisms that do not restrict the input format have large distortion. Our lower bound is nearly tight, since there are truthful ordinal mechanisms that have distortion $O(\sqrt{m \log m})$, implying that the loss from restricting the input format does not impose a significant additional burden.

The gap between the upper and lower bounds ($O(\sqrt{m \log m})$ and $\Omega(\sqrt{m})$) remains open. However, the seminal work by Gibbard [16] shows that an ordinal mechanism is truthful iff it is localized and non-perverse. An interesting question is if there exist mechanisms which approach the distortion lower bound, and violate these properties, extended to cardinal mechanisms. We show that indeed such mechanisms exist, and give one such mechanism with distortion $O(\sqrt{m \log m})$, matching the best known upper bound.

We will use the following definition and characterization for our proofs. Recall that we assume that for each voter, no two alternatives have the same utility. Since we assume truthfulness in this section, this extends to their strategies as well.

► **Definition 4.** A mechanism is a dictatorship if there exists voter i so that for any strategy profile \vec{u} , $\mu(\vec{u}) = \arg \max_{a \in A} u_i(a)$. Voter i is said to decide the mechanism in this case.

⁴ A mechanism that is Pareto-optimal must be unanimous, hence the lower bound holds for all Pareto-optimal mechanisms as well.

► **Definition 5.** Cardinal mechanism μ is unanimous if whenever there exists an alternative $a^* \in A$ such that $\arg \max_{a \in A} u_i(a) = a^*$ for all voters i , $\mu(\vec{u})$ selects a^* with probability 1.

► **Theorem 6** ([18, 10]). *A unanimous truthful cardinal mechanism is a randomization over dictatorial mechanisms.*

A similar result was also shown for ordinal mechanisms by Gibbard [16]. We first show tight bounds for deterministic mechanisms.

► **Theorem 7.** *Deterministic truthful cardinal mechanisms have distortion $\Theta(mn)$.*

Proof. For the upper bound, consider the mechanism that picks the maximum utility alternative for voter 1. This mechanism has welfare at least $1/m$, while the maximum welfare obtainable is n , giving us the upper bound.

For the lower bound, we first assume that for every $a \in A$, there is some strategy profile for which the mechanism returns a . This is without loss of generality, since if there is some alternative a that is never chosen, then when all agents have utility 1 for a and 0 for all other alternatives, the distortion is infinite. We next show in the following claim that a truthful deterministic cardinal mechanism must be unanimous.

► **Claim 8.** *A truthful deterministic cardinal mechanism must be unanimous.*

Proof. Suppose the truthful deterministic mechanism μ is not unanimous. Then for some alternative a , \vec{u} is a utility profile where the maximum utility alternative for every voter is a , but the mechanism returns $b \neq a$. Since a is feasible, there exists \vec{u}' such that $\mu(\vec{u}') = a$. Let the voters deviate, one by one, from their strategy in \vec{u} to \vec{u}' . For some voter i , the mechanism chooses a after the player deviates, and not before. This player then has an incentive to report her utility as in \vec{u}' when her actual utility is as in \vec{u} , when the other voters have utilities as in the utility profile before (and also after) the deviation by i . The mechanism thus cannot be truthful. ◀

We can now complete the proof of the theorem. By Theorem 6 and Claim 8, any truthful deterministic cardinal mechanism must be a dictatorship. For such a mechanism, let i be the voter that decides the mechanism. Consider the utility profile where voter i has utility $1/m + \epsilon$ for candidate a and utility $1/m - \epsilon/(m - 1)$ for the other alternatives. All the other voters have utility 1 for some candidate $b \neq a$. Then the maximum welfare is about $(n - 1)$ while the mechanism obtains welfare $1/m + \epsilon$, giving us the required distortion. ◀

► **Theorem 9.** *Unanimous truthful cardinal mechanisms have distortion $\Omega(m)$.*

Proof. Let the unanimous mechanism be μ , by Theorem 6 this must be a randomization over dictatorships. Select $a^* \in A$, and consider the utility profile \vec{u} where all agents have utility $0.5 - \epsilon$ for a^* and $0.5 + \epsilon$ for some other alternative uniformly selected from $A \setminus \{a^*\}$. Then $\mu(\vec{u})$ selects a^* with probability 0, and gets expected welfare $n/2(m - 1)$, while alternative a^* has welfare $n/2$ (we assume $\epsilon \rightarrow 0$), giving distortion $\Omega(m)$. ◀

We now show a lower bound for all truthful cardinal mechanisms.

► **Theorem 10.** *Any truthful cardinal mechanism has distortion $\Omega(\sqrt{m})$.*

Proof. We will assume that \sqrt{m} is an integer and n is divisible by \sqrt{m} . This helps simplify the proof but is not required for it to hold. Let μ be a truthful mechanism. Let $\{a_1, a_2, \dots, a_{\sqrt{m}}\} = A^* \subseteq A$ be a subset of alternatives of size \sqrt{m} . Partition the set of agents N into \sqrt{m} sets of equal size n/\sqrt{m} , say $N_1, N_2, \dots, N_{\sqrt{m}}$. Let $n' := n - (n/\sqrt{m})$.

Create a utility profile where for each $i \in \{1, 2, \dots, \sqrt{m}\}$, the agents in set N_i have utility 1 for a_i and utility 0 for the other alternatives. Call this profile \vec{u}^0 . For this profile, at least one of the alternatives in A^* is selected by μ with probability $\leq 1/\sqrt{m}$, say alternative $a_{\sqrt{m}}$. Let the probability it gets be $p_0 \leq 1/\sqrt{m}$.

Now, for all voters $i \in N \setminus N_{\sqrt{m}} = \{1, \dots, n'\}$ we will make the utility uniform among the alternatives in $A \setminus \{a_{\sqrt{m}}\}$, one voter at a time. Formally, let v be a utility function where $a_{\sqrt{m}}$ has utility 0 and all other alternatives have utility $1/(m-1)$. Let utility profile $\vec{u}^i = \vec{u}^{i-1}/_i v$, and let p_i be the probability that alternative $a_{\sqrt{m}}$ gets for profile \vec{u}^i , for $i \in \{1, 2, \dots, n'\}$.

Observe that only the utility for the i th agent changes when we change the strategy profile from \vec{u}^{i-1} to \vec{u}^i . Suppose voter i has utility function v . Then the truthfulness condition requires that the probability mass on alternatives other than $a_{\sqrt{m}}$ should be maximum when it reports truthfully, and hence $1 - p_i \geq 1 - p_{i-1}$, or $p_i \leq p_{i-1}$. Thus, $p_{n'} \leq p_0 \leq 1/\sqrt{m}$. Now for the bound on the distortion, for the last utility profile $\vec{u}^{n'}$ when only voters in $N_{\sqrt{m}}$ have positive utility 1 for $a_{\sqrt{m}}$ and all other voters divide their utility equally among the other alternatives, $a_{\sqrt{m}}$ has maximum social welfare equal to n/\sqrt{m} , and this is picked with probability at most $1/\sqrt{m}$. The distortion bound follows from simple calculations. ◀

Randomized Hyperspherical Mechanisms

We now describe a truthful cardinal mechanism that has distortion $O(\sqrt{m \log m})$, matching the best known distortion upper bound, and which violates the properties of localization and non-perversity. We first present the mechanism and its analysis, and then show an example for which the mechanism violates these properties. For a dimension m , let $\mathbf{1}$ be the all-ones vector. The standard $(m-1)$ -simplex $\{x \in \mathbb{R}_{\geq 0}^m : \|x\|_1 = 1\}$ is denoted Δ_m . Before we describe our mechanism, for a fixed radius $R \geq 0$ and dimension m , consider the following sets:

$$S_m^1(R) = \left\{ p \in \mathbb{R}^m : \left\| p - \frac{1}{m} \mathbf{1} \right\|_2 \leq R, \|p\|_1 = 1 \right\}, \quad S_m^2(R) = S_m^1 \cap \mathbb{R}_{\geq 0}^m.$$

$S_m^1(R)$ is the set of points in \mathbb{R}^m whose coordinates (possibly negative) sum to 1, and are at distance at most R from $(1/m, \dots, 1/m)$. $S_m^2(R)$ is the set of points that lie in the intersection of the standard simplex with the ball of radius R with center $\frac{1}{m} \mathbf{1}$. Note that both sets are convex. Given $x \in \mathbb{R}^m$, there is a boundary point p that maximizes $p^T x$ in either of these sets, since the objective is linear.

We now describe our mechanism. Let \vec{u} be the given strategy profile. Let μ_1 be the mechanism that picks an alternative with uniform probability and returns it. Let μ_2 be the mechanism that selects a radius R uniformly from the set

$$\Gamma = \left\{ \frac{1}{\sqrt{m(m-1)}}, \frac{2}{\sqrt{m(m-1)}}, \frac{4}{\sqrt{m(m-1)}}, \dots, \frac{m-1}{\sqrt{m(m-1)}} \right\}$$

and selects a voter i from N with uniform probability. Mechanism μ_2 returns the point $p \in S_m^2(R)$ that maximizes $p^T u_i$. Since p lies in the standard simplex, it is a distribution. Finally, our randomized hyperspherical mechanism runs μ_1 with probability $1/2$, and μ_2 with probability $1/2$.

Analysis. The truthfulness of the mechanism is evident since for μ_2 , the voter i and radius R are chosen independently from the input \vec{u} , and we choose $p \in S_m^2(R)$ that maximizes the expected utility $p^T u_i$ for voter i . There are thus two things we need to show: that the

Algorithm 1 Compute- p .

Input: Dimension $m \in \mathbb{Z}_+$, vector $x \in \Delta_m$, radius $R \geq 0$.

Output: Distribution $p \in S_m^2(R)$ that maximizes $p^T x$.

- 1: $p \leftarrow \frac{1}{m} \mathbf{1} + R \frac{x - \frac{1}{m} \mathbf{1}}{\|x - \frac{1}{m} \mathbf{1}\|_2}$
 - 2: **if** $p \geq 0$ **then**
 - 3: **return** p
 - 4: **else**
 - 5: $x' \leftarrow \frac{x[1:m-1]}{\sum_{i=1}^{m-1} x_i}$, $R' \leftarrow \sqrt{R^2 - \frac{1}{m(m-1)}}$, $m' \leftarrow m - 1$.
 - 6: $p' \leftarrow \text{Compute-}p(m', x', R')$. **return** $(p' \ 0)$.
-

distribution p that maximizes $p^T u_i$ over $S_m^2(R)$ can be obtained efficiently, and that the mechanism has distortion $O(\sqrt{m \log m})$.

We first give an algorithm for computing p . Let $x = u_i$ be the utility for the voter chosen. We reindex the alternatives so that $x_1 \geq \dots \geq x_m$. We use $x[1:k] = (x_1, \dots, x_k)$ to denote the vector consisting of the first k components of x .

Theorem 11 shows that the algorithm finds the point $p \in S_m^2(R)$ that maximizes $p^T x$, as required. The algorithm first finds the point in $p \in S_m^1(R)$ that maximizes $p^T x$. If p is nonnegative, then $p \in S_m^2(R)$, and this is returned. If not, then we show that in the optimal distribution, it must be the case that $p_m = 0$. In this case, let x' and R' be x and R as modified in Line 5. It can be checked that $S_{m-1}^2(R')$ is the intersection of $S_m^2(R)$ with the hyperplane $p_m = 0$. In this case, we focus on the first $m - 1$ coordinates of x , and recursively find the point in p' in $S_{m-1}^2(R')$ that maximizes $p'^T x'$. We show in the proof that the distribution $p = (p' \ 0)$ is the point in $S_m^2(R)$ that maximizes $p^T x$.

► **Theorem 11.** *Algorithm 1 correctly returns $p \in S_m^2(R)$ that maximizes $p^T x$.*

Proof. The proof of the theorem follows from these claims.

► **Claim 12.** *The point p obtained in Line 1 is the point $q \in S_m^1(R)$ that maximizes $q^T x$.*

Proof. It can be checked from the steps in the algorithm that $\sum_i p_i = 1$, and $\|p - \mathbf{1}/m\|_2 = R$, hence $p \in S_m^1(R)$. Secondly, $q \in S_m^1(R)$ maximizes $q^T x$ iff q maximizes $(q - \mathbf{1}/m)^T x = \|q - \mathbf{1}/m\|_2 \|x\|_2 \cos \theta$, where θ is the angle between $q - \mathbf{1}/m$ and x . The point p is chosen such that $\theta = 0$ and $\|q - \mathbf{1}/m\|_2 = R$, hence it maximizes $q^T x$. ◀

Clearly, if $p \geq 0$, then $p \in S_m^2(R)$ and we are done. Else, since $x_1 > x_2 > \dots > x_m$, $p_1 \geq p_2 \geq \dots \geq p_m$ (else permuting the coordinates of p to obtain these inequalities would give us a point in $S_m^1(R)$ with higher value for $p^T x$). Hence suppose $p_m < 0$.

► **Claim 13.** *If $p_m < 0$, then there is a point $p' \in S_{m-1}^2(R')$ that maximizes $q'^T x'$ over all points q in $S_m^2(R)$, and has $p'_m = 0$.*

Proof. Let q' be a point in $S_m^2(R)$ that maximizes $q'^T x'$ over such points, and suppose $q'_m > 0$. Then $x^T q' \leq x^T p$, and since $q'_m > 0$, $p_m < 0$, and the coordinates for each of these vectors is nonincreasing in the indices, there is a point p' on the line joining q' and p that lies in $S_m^1(R)$ (since both these points lie in $S_m^1(R)$, and this set is convex) so that $x^T p' \geq x^T q'$ and $p'_m = 0$, and with coordinates nonincreasing in the indices. Hence $p' \geq 0$, and hence $p' \in S_m^2(R)$, which is the required point. ◀

Now let $\lambda = \sum_{i=1}^{m-1} x_i$, $x' = (x_1, x_2, \dots, x_{m-1})/\lambda$, and $R' = \sqrt{R^2 - \frac{1}{m(m-1)}}$, as we use in Line 5. Let q' be a point in $S_{m-1}^2(R')$ that maximizes $q'^T x[1 : m-1]$ over all such points q .

► **Claim 14.** *If $p_m < 0$, then $(q' 0) \in S_m^2(R)$ maximizes $q^T x$ over all such points q .*

Proof. Since $q' \in S_{m-1}^2(R')$, $(q' 0)$ is in $S_m^2(R)$. For the second part, let p' be as obtained in Claim 13, and let $q'' = (p'_1, \dots, p'_{m-1})$. Since $p'_m = 0$, $q'' \in S_{m-1}^2(R')$. Then $\lambda q'^T x' = (q' 0)^T x$, and similarly $\lambda q''^T x' = (q'' 0)^T x = p'^T x$. Suppose for a contradiction that $(q' 0)^T x = \lambda q'^T x' < p'^T x$. Then $\lambda q'^T x' < \lambda q''^T x'$, which contradicts the optimality of q' . ◀

Thus, if the point $p \geq 0$, then this is a point in $S_m^2(R)$ that maximizes $p^T x$, and is correctly returned. If not, then $p_m < 0$. In this case, by Claim 13, $p'_m = 0$, and by Claim 14 it is sufficient to compute the point $q' \in S_{m-1}^2(R')$ that maximizes $q^T x$ over all such points q , and return the vector $(q' 0)$, which is the iterative step in our algorithm as well. ◀

We now show the bound on the distortion.

► **Theorem 15.** *The randomized hyperspherical mechanism has distortion $O(\sqrt{m \log m})$.*

Proof. Let $a^* \in A$ be the alternative with maximum social welfare. Observe that mechanism μ_1 , that picks an alternative with uniform probability, has expected welfare of n/m . If $\text{sw}(a^*) \leq n\sqrt{\frac{\log m}{m}}$, then since μ_1 is picked w.p. $1/2$, $\text{dist}(\mu) \leq 2 \text{dist}(\mu_1) \leq 2\sqrt{m \log m}$.

Else, assume μ_2 is the mechanism picked. Let $p(a)$ be the probability that alternative a is picked, and $p_i(a)$ be the probability that agent i and alternative a are picked. Then $p(a) = \sum_{i \in N} p_i(a)$. For any $i \in N$ and corresponding utility function u_i , there is hypersphere with radius R between $\|u_i - 1/m\|_2$ and $\|u_i - 1/m\|/2$. The point p on this hypersphere that maximizes $p^T u_i$ is the point on the line joining $1/m$ with u_i , which clearly lies in the simplex. Since this point is at least halfway to u_i , the coordinate corresponding to alternative a has value $\lambda u_i(a) + (1 - \lambda)\frac{1}{m}$ for $\lambda \geq 1/2$, and hence


$$p_i(a) \geq \frac{1}{n} \frac{1}{\log m} \frac{u_i(a)}{2}$$

and hence $p(a) \geq \sum_i p_i(a) \geq \text{sw}(a)/(2n \log m)$. Since $\text{sw}(a^*) \geq n\sqrt{\frac{\log m}{m}}$, the distortion is

$$\text{dist}(\mu) \leq 2 \text{dist}(\mu_2) \leq \frac{2\text{sw}(a^*)}{\text{sw}(a^*)p(a^*)} \leq \frac{4n \log m}{\text{sw}(a^*)} \leq \frac{4n \log m}{n\sqrt{\log m/m}} = 4\sqrt{m \log m}. \quad \blacktriangleleft$$

We now show that the randomized hyperspherical mechanism violates the properties of non-perversity and locality. Let there be 3 alternatives and 1 voter. The mechanism randomizes over hyperspheres with radii $1/\sqrt{6}$ and $2/\sqrt{6}$ with probability $1/4$ each, and selects an alternative uniformly with probability $1/2$.

- Perverse. Consider the utility profiles $u = (1/4 - \epsilon, 1/2, 1/4 + \epsilon)$ and $u' = (0, 3/4, 1/4)$, where $\epsilon \rightarrow 0^+$. The relative ordering of the alternatives for both the profiles is same. Running the mechanism returns the distributions $p \approx (0.208, 0.584, 0.208)$ and $p' \approx (0.187, 0.579, 0.234)$ for u and u' , respectively. The mechanism is perverse, since the probability of the second alternative decreases, despite its utility increasing and the ordering of the alternatives remaining unchanged.



■ **Figure 1** The hyperspherical mechanism for a single sphere of radius $1/\sqrt{6}$, with three alternatives.

- Not localized. Consider the utility profiles $u = (1 - \epsilon, \epsilon, 0)$ and $u' = (3/4, 1/4, 0)$, where $\epsilon \rightarrow 0^+$. The total utility for the first two alternatives is same for both profiles and is equal to 1. Running the mechanism returns the distributions $p \approx (0.584, 0.208, 0.208)$ and $p' \approx (0.579, 0.234, 0.187)$ for u and u' respectively. The mechanism is not localized, since the total probability given to the first two alternatives is different: 0.792 and 0.813, despite having the same total utility.

Figure 1 shows the projection of the two utility vectors in the first example onto the second hypersphere of radius $1/\sqrt{6}$.

4 Approximately Truthful Mechanisms

Given the strong lower bounds on distortion with truthful mechanisms, we now consider approximately truthful mechanisms. A mechanism is δ -truthful if no voter can increase her expected utility by more than an additive δ by misreporting her utilities. In this case, perhaps surprisingly, we are able to show mechanisms that obtain near-optimal distortion. Our mechanism takes a parameter $\delta \in (0, 1)$ as input. The resulting mechanism is 2δ -truthful, and has distortion that goes to 1 as the number of voters increases. In particular, if $\delta = 25m/n < 1$, then the mechanism has distortion almost 1.01.

The mechanism proceeds as follows. It first elicits the strategy profile \vec{u} of the voters. For alternative $j \in [m]$, define $s_j = \text{sw}(j) = \sum_{i \in N} u_i(j)$. Assume (or re-index the set of alternatives) that $s_1 \geq s_2 \geq \dots \geq s_m$. Define

$$\lambda := \max \left\{ k \in [m] : \sum_{i=1}^k (s_i - s_k) < \frac{1}{\delta} \right\}. \quad (1)$$

Note that $\lambda \geq 1$. Then the mechanism returns the probability distribution defined as follows:

$$p_k = \begin{cases} \frac{1}{\lambda} \left(1 - \delta \sum_{i=1}^{\lambda} (s_i - s_k) \right) & \text{for } k \leq \lambda \\ 0 & \text{for } k > \lambda \end{cases} \quad (2)$$

Since $\sum_{k=1}^{\lambda} \sum_{i=1}^{\lambda} (s_i - s_k) = 0$, the sum of probabilities $\sum_{i=1}^m p_i = \sum_{i=1}^{\lambda} p_i = 1$. Further,

for all $k \leq \lambda$ the sum

$$\sum_{i=1}^{\lambda} (s_i - s_k) = \sum_{i=1}^k (s_i - s_k) + \sum_{i=k+1}^{\lambda} (s_i - s_k) \leq \sum_{i=1}^k (s_i - s_k) \leq 1/\delta,$$

where the first inequality is by the indexing of the alternatives and the second is by definition of λ . Hence all the probabilities are non-negative. It can also be shown by a quick calculation that the probability distribution returned by the mechanism can also be written as:

$$p_k = \begin{cases} p_1 - \delta(s_1 - s_k) & \text{for } k \leq \lambda \\ 0 & \text{for } k > \lambda \end{cases} \quad (3)$$

with p_1 chosen so that the sum of the probabilities is 1. We now show the required properties for this mechanism.

► **Theorem 16.** *Given parameter $\delta \in (0, 1)$, the described mechanism has distortion $< \frac{s_1}{s_1 - \frac{1}{4\delta}}$.*

Proof. Clearly, the maximum welfare obtainable is s_1 . The expected welfare obtained by the mechanism is $\sum_{j=1}^m p_j s_j$, and replacing for s_j from (3), this gives us $s_1 - (p_1 - \sum_{j=1}^{\lambda} p_j^2)/\delta$. Optimizing over the p_j 's, we find that the worst social welfare is obtained when $p_1 = (1 + 1/\lambda)/2$, $p_j = 1/(2\lambda)$ for $2 \leq j \leq \lambda$, and the expected welfare in this case is at least $s_1 - (1/4\delta)$. The bound on the distortion follows immediately. ◀

► **Theorem 17.** *Given parameter $\delta \in (0, 1)$, the described mechanism is 2δ -truthful.*

Proof. We show the following stronger property: give two strategy profiles \vec{u}, \vec{u}' , let $\vec{s} := (\sum_{i \in N} u_i(a))_{a \in A}$ and $\vec{s}' := (\sum_{i \in N} u'_i(a))_{a \in A}$ be the respective social welfare vectors. Let $\alpha := \|\vec{s} - \vec{s}'\|_1$ be the L_1 distance between the two welfare vectors. Then the distributions returned by the mechanism given inputs \vec{u} and \vec{u}' differs in any component by at most $\alpha\delta$. The theorem then follows, since by deviating, a single player can change the total welfare by at most 2, and hence the distribution changes in any component by at most 2δ . We first state the following claim, which states that if λ remains unchanged for \vec{s} and \vec{s}' , then the property described holds. Let p and p' be the probability distributions returned by the mechanism for \vec{u} and \vec{u}' (with welfare vectors \vec{s} and \vec{s}' respectively). The proof follows immediately from (2).

► **Claim 18.** *If λ is the same for \vec{s} and \vec{s}' , then the distributions p and p' differ in each component by at most $\delta\|\vec{s} - \vec{s}'\|_1$.*

The next claim shows that if there exists an index k so that $\sum_{i=1}^k (s_i - s_k) = \frac{1}{\delta}$, then including this in λ does not change the distribution. Another way of viewing the claim is that it shows that the probability distribution changes continuously with s . In particular, the strict inequality in the definition of λ can be replaced by an inequality without changing the distribution.

► **Claim 19.** *Given a strategy profile \vec{u} , let $\vec{s} := \sum_{i \in N} u_i$. Let λ be as defined previously, and define λ' as any index k so that $\sum_{i=1}^k (s_i - s_k) = \frac{1}{\delta}$ (if it exists). Then the distribution p is unchanged if we replace λ by λ' in (2).*

Proof. Assume λ' exists, else the claim is trivially true. Let p be as defined in (2), and p' be the distribution obtained from (2) with λ replaced by λ' . By definition, $\lambda' > \lambda$. Let $r := \lambda' - \lambda$. Then since $1/\delta = \sum_{i=1}^{\lambda+k} (s_i - s_{\lambda+k})$ for all $k \in \{1, \dots, r\}$, it must be true that

27:14 On the Welfare of Cardinal Voting Mechanisms

$s_{\lambda+1} = \dots = s_{\lambda'}$. Hence, it must also be true that $\sum_{i=1}^{\lambda} (s_i - s_{\lambda'}) = \sum_{i=1}^{\lambda'} (s_i - s_{\lambda'})$. We will use this last equality later in the proof. For any $k \leq \lambda'$, we get the probability distribution

$$\begin{aligned} p'_k &= \frac{1}{\lambda'} \left(1 - \delta \sum_{i=1}^{\lambda'} (s_i - s_k) \right) \\ &= \frac{1}{\lambda'} \left(1 - \delta \sum_{i=1}^{\lambda'} (s_i - s_{\lambda'}) - \delta \lambda' (s_{\lambda'} - s_k) \right) \\ &= \frac{1}{\lambda'} (\delta \lambda' (s_k - s_{\lambda'})) = \delta (s_k - s_{\lambda'}) \end{aligned}$$

Note that (i) for $\lambda < k \leq \lambda'$, since $s_k = s_{\lambda'}$ as discussed above, $p'_k = 0 = p_k$, and similarly for $k > \lambda'$, by definition $p'_k = 0 = p_k$. Hence it remains to show that for $k \leq \lambda$, $p_k = \delta (s_k - s_{\lambda'})$. Simple calculations yield that for $k \leq \lambda$,

$$p_k - \delta (s_k - s_{\lambda'}) = \frac{1}{\lambda} \left(1 - \delta \sum_{i=1}^{\lambda} (s_i - s_{\lambda'}) \right) = 0,$$

where the last equality follows from the discussion in the first paragraph. \blacktriangleleft

We now complete the proof of the theorem. Instead of the strategy profiles, we consider directly the resulting welfare vectors. Consider the straight line from \vec{s} to \vec{s}' . Let λ and λ' be defined as in (1) for s and s' respectively. If $\lambda = \lambda'$, then by Claim 18, the theorem holds. Suppose instead that $\lambda' = \lambda + r$, for some $r \geq 1$. Let $\vec{s}^0 = \vec{s}$ and $\vec{s}^r = \vec{s}'$. We segment the path from \vec{s} to \vec{s}' into r segments $[\vec{s}^0, \vec{s}^1]$, $[\vec{s}^1, \vec{s}^2]$, \dots , $[\vec{s}^{r-1}, \vec{s}^r]$ so that (i) λ remains unchanged at each point within a segment, and (ii) at the i th breakpoint \vec{s}^i , $1/\delta = \sum_{j=1}^{\lambda+i} (s_j - s_{\lambda+i})$. It follows from the previous claims that for any component, the change in probability between \vec{s} and \vec{s}' is at most $\delta \sum_{i=1}^r \|\vec{s}^i - \vec{s}^{i-1}\|_1 = \delta \|\vec{s}^r - \vec{s}\|_1$, as required to complete the proof. \blacktriangleleft

5 Convergence and Price of Anarchy in Iterative Voting

In this section, we focus on the deterministic mechanism μ that given a strategy profile \vec{u}' , chooses the alternative that maximizes the reported social welfare $\sum_{i \in N} u'_i(a)$. The alternatives are indexed, and ties are resolved in favour of the alternative with lower index. All results are presented for this mechanism. It is easily seen that even for two voters, this mechanism is not truthful. Hence we focus on the stable outcomes of strategic voting, in particular strategy profiles that are at a pure Nash equilibrium. As before, we are concerned with the social welfare of the alternative chosen by the mechanism at equilibrium. The distortion in this case is equivalent to the Price of Anarchy, and we refer to it as such here. The PoA is in general unbounded, and so are certain refinements. Instead, we consider equilibria which arise as a result of iterative voting dynamics, when starting from the initial truthful utility profile, a voter deviates at each step in a manner that improves her utility, until the voters reach an equilibrium. In contrast to truthful mechanisms, we show a strong positive result. We show that a particular and natural class of iterative voting dynamics converges quickly to an equilibrium. Further, the price of anarchy for the class of equilibria thus obtained is close to 1, as the number of voters increases. We note that while previous papers have studied either the convergence of iterative voting (e.g., [19, 22, 25]) or the PoA obtained for mechanisms such as plurality and veto (e.g., [9]), ours is the first to obtain results for the PoA of outcomes obtained by iterative voting for a natural cardinal mechanism.

The PoA over all equilibria is unbounded; consider a simple example where all voters have utility 1 for candidate a and 0 for candidate b , but all choose the strategy with utility 1 for b . This is clearly an odd example, and there seem to be two possible remedies. Firstly, we could consider strong equilibria, where \vec{u} is a strong equilibrium if no set (or coalition) of players can deviate to improve the expected utility of all the players in the set. Secondly, we could consider truth-biased agents, who prefer to vote truthfully if it does not reduce their utility (for which positive results are known in some cases, see [20, 23]). Unfortunately, in both cases we show that the PoA continues to be unbounded.

► **Theorem 20.** *The PoA with truth-biased agents is unbounded. Further, an equilibrium may not exist.*

► **Theorem 21.** *The PoA of strong Nash equilibria is unbounded.*

Given these negative results, we focus on equilibrium outcomes that are obtained as a result of iterative voting. As stated, we fix the mechanism μ that given a strategy profile \vec{u} , chooses the alternative that maximizes the reported social welfare $\sum_{i \in N} u'_i(a)$. The alternatives are indexed, and ties are resolved in favour of the alternative with lower index. We assume that initially, all voters report their utilities truthfully. At each step, a single voter chooses a different strategy that improves her utility. We say that a particular iterative voting dynamics converges if in finite time, the strategy profile is an equilibrium. We are interested in the PoA of equilibria that are obtained as a result of iterative voting.

Again, we show in the appendix that without further restrictions, the PoA for equilibria obtained can be unbounded, even if the deviating player at each step strictly improves her utility. Let us instead consider the iterative voting process where the deviation by the player at each step satisfies the following properties:

- (A) The utility of the deviating player must strictly increase after the deviation.
- (B) The deviating player can increase the reported utility for a single alternative, and this alternative must be chosen by the mechanism after the deviation.

With these restrictions, the PoA for the class of equilibria obtained is nearly 1.

► **Theorem 22.** *The PoA for iterative voting with restrictions (A) and (B) is $\frac{\max_{a \in A} \text{sw}(a)}{\max_{a \in A} \text{sw}(a) - 2 \log_2 m}$.*

Proof. Let \vec{v}^t be the strategy profile in the t th time step. Then $\vec{v}^0 = \vec{u}$, where \vec{u} is the utility profile for the voters. We define $\text{sw}^t(a) := \sum_{i \in N} u_i^t(a)$ as the welfare of alternative a according to the strategy profile at step t . Then $\text{sw}^0(a) = \text{sw}(a)$ since iterative voting starts with the true utility as strategy, and we index the alternatives so that $\text{sw}(a_1) \geq \text{sw}(a_2) \geq \dots \geq \text{sw}(a_m)$. In particular, the maximum-welfare candidate is a_1 . We say an alternative wins at time t if it maximizes $\text{sw}^t(a)$, and among all such alternatives, has the lowest index.

Fix any $j \in \{2, \dots, m\}$, and let t be the first time that an alternative a_k with $k \geq j$ wins, hence $\text{sw}(a_j) \geq \text{sw}(a_k)$ by our indexing. Further, since this is the first time that a_k wins, it is also the first time that any voter raises its utility for a_k , and hence $\text{sw}(a_k) \geq \text{sw}^{t-1}(a_k)$. Lastly, since the voter that deviates at time t changes its utility by at most 1 for any alternative, and a_k wins at time t , it must be true that $\text{sw}^{t-1}(a_k) \geq \max_{r \leq m} \text{sw}^{t-1}(a_r) - 2$. Putting these together,

$$\text{sw}(a_j) \geq \text{sw}(a_k) \geq \text{sw}^{t-1}(a_k) \geq \max_{r \leq m} \text{sw}^{t-1}(a_r) - 2 \geq \frac{1}{j-1} \sum_{r=1}^{j-1} \text{sw}^{t-1}(a_r) - 2$$

where the last inequality is simply because the maximum of set of numbers is at least its average. Now observe that step t is the first step when an alternative with index at least j

has won, and hence by the dynamics restriction, this is the first step when an alternative with index at least j had its utility increased. Hence $\sum_{r \geq j} \text{sw}^{t-1}(a_r) \leq \sum_{r \geq j} \text{sw}(a_r)$. Since the sum of utilities of a strategy profile is always n , $\sum_{r < j} \text{sw}^{t-1}(a_r) \geq \sum_{r < j} \text{sw}(a_r)$. Plugging this into the previous inequality yields

$$\text{sw}(a_j) \geq \frac{1}{j-1} \sum_{r=1}^{j-1} \text{sw}(a_r) - 2$$

Let a_j be the highest indexed alternative to win during the dynamics. Then the above inequality is valid for all $k \in \{2, \dots, j\}$, giving us a recurrence relation. To solve this recurrence, we can check that the hypothesis that $\text{sw}(a_j) \geq \text{sw}(a_1) - 2 \log_2 j$ is correct. Our proof thus shows the stronger property that if alternative a wins at any time step in the dynamics, then $\text{sw}(a) \geq \text{sw}(a_1) - 2 \log_2 m$. ◀

Unfortunately, it turns out that iterative voting even with these restrictions may not converge (Theorem 24, in the Appendix). However, with one further restriction on the allowable deviations, we can prove convergence.

(C) The deviating player can decrease the reported utility for a single alternative, and this alternative must be chosen by the mechanism before the deviation.

It is not hard to see that the number of steps required for convergence depends upon the least value by which a voter can change her score. As an example, consider two alternatives and two voters with utilities $(0.5 - \epsilon, 0.5 + \epsilon)$ and $(0.5 + \epsilon, 0.5 - \epsilon)$. Let δ be the least value by which a voter can change her utility. Then each time a voter increases the reported utility of her preferred alternative by δ , the alternative chosen by the mechanism changes, and hence convergence takes $\Omega(1/\delta)$ steps. In fact, a convergence bound of $O(mn/\delta)$ for iterative voting with restrictions (A), (B) and (C) can easily be shown, by observing that in each move a voter shifts her stated utility from a less preferred alternative to a more preferred alternative by at least δ . A voter can thus move at most $\frac{m}{\delta}$ times, and hence after $\frac{mn}{\delta}$ steps the iterative voting process must reach a Nash equilibrium.

We can obtain even better convergence bounds for iterative voting, where apart from the initial votes, in all subsequent strategies of a voter, exactly one alternative is given utility 1. These subsequent votes are then plurality votes, for which Meir et al. [20] show a bound of $O(mn)$ on the convergence. Hence every $O(mn)$ steps a new voter must change to plurality voting from her initial utility, and hence this iterative voting process must converge in $O(mn^2)$ steps. If the dynamics also satisfies restrictions (A) and (B), then the equilibrium obtained has PoA as shown in Theorem 22.

References

- 1 Salvador Barbera. Nice decision schemes. In *Decision theory and social ethics*, pages 101–117. Springer, 1978.
- 2 Salvador Barbera. An introduction to strategy-proof social choice functions. *Soc Choice Welfare* 18, pages 619–653, 2001.
- 3 Salvador Barbera, Anna Bogomolnaia, and Hans Van Der Stel. Strategy-proof probabilistic rules for expected utility maximizers. *Mathematical Social Sciences*, 35(2):89–103, 1998.
- 4 Gerdus Benade, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Preference Elicitation For Participatory Budgeting. In *Association for Advancement of Artificial Intelligence (AAAI), February 4 - 9, 2017, San Francisco, California, USA*, 2017. Forthcoming.

- 5 Umang Bhaskar, Varsha Dani, and Abheek Ghosh. Truthful and Near-Optimal Mechanisms for Welfare Maximization in Multi-Winner Elections. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*, 2018.
- 6 Eleanor Birrell and Rafael Pass. Approximately strategy-proof voting. In *IJCAI*, pages 67–72, 2011.
- 7 Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artif. Intell.*, 227:190–213, 2015. doi:10.1016/j.artint.2015.06.003.
- 8 Felix Brandt, Vincent Conitzer, Ulle Endriss, Ariel D Procaccia, and Jérôme Lang. *Handbook of computational social choice*. Cambridge University Press, 2016.
- 9 Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D Procaccia. How bad is selfish voting? In *AAAI*, volume 13, pages 138–144, 2013.
- 10 Bhaskar Dutta, Hans Peters, and Arunava Sen. Strategy-proof cardinal decision schemes. *Social Choice Welfare*, 28:163–179, 2007.
- 11 Edith Elkind, Evangelos Markakis, Svetlana Obraztsova, and Piotr Skowron. Equilibria of plurality voting: Lazy and truth-biased voters. In *International Symposium on Algorithmic Game Theory*, pages 110–122. Springer, 2015.
- 12 Uriel Feige and Moshe Tennenholtz. Responsive Lotteries. In *Algorithmic Game Theory - Third International Symposium, SAGT 2010, Athens, Greece, October 18-20, 2010. Proceedings*, pages 150–161, 2010.
- 13 Aris Filos-Ratsikas and Peter Bro Miltersen. Truthful approximations to range voting. *International Conference on Web and Internet Economics*, pages 175–188, 2014.
- 14 Sumit Ghosh, Manisha Mundhe, Karina Hernandez, and Sandip Sen. Voting for movies: the anatomy of a recommender system. In *Proceedings of the third annual conference on Autonomous Agents*, pages 434–435. ACM, 1999.
- 15 Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.
- 16 Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica: Journal of the Econometric Society*, pages 665–681, 1977.
- 17 Claude Hillinger. The case for utilitarian voting. *SSRN*, 2005. doi:10.2139/ssrn.732285.
- 18 Aanund Hylland. Strategy proofness of voting procedures with lotteries as outcomes and infinite sets of strategies. *Unpublished paper, University of Oslo.[341, 349]*, 1980.
- 19 Omer Lev and Jeffrey S Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 611–618. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- 20 Reshef Meir, Maria Polukarov, Jeffrey S Rosenschein, and Nicholas R Jennings. Convergence to Equilibria in Plurality Voting. In *AAAI*, volume 10, pages 823–828, 2010.
- 21 Shasikanta Nandeibam. An alternative proof of Gibbard’s random dictatorship result. *Social Choice and Welfare*, 15(4):509–519, 1998.
- 22 Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the Convergence of Iterative Voting: How Restrictive Should Restricted Dynamics Be? In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 993–999, 2015.
- 23 Svetlana Obraztsova, Evangelos Markakis, and David RM Thompson. Plurality voting with truth-biased agents. In *International Symposium on Algorithmic Game Theory*, pages 26–37. Springer, 2013.

- 24 Ariel D Procaccia and Jeffrey S Rosenschein. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*, pages 317–331. Springer, 2006.
- 25 Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey S. Rosenschein. Analysis of Equilibria in Iterative Voting Schemes. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 1007–1013, 2015.
- 26 Reyhaneh Reyhani and Mark Wilson. Best-reply dynamics for scoring rules. In *20th European Conference on Artificial Intelligence*. IOS Press, 2012.
- 27 Mark Allen Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, 10(2):187–217, 1975.

6 Appendix

In the appendix, we give the missing proofs from Section 5. We first prove that the PoA with truth-biased voters is unbounded, and in fact a pure Nash equilibrium may not exist.

Proof of Theorem 20. We will give an example for which the price of anarchy $\rightarrow \infty$. There are 3 alternatives $\{a, b, c\}$ and $2n + 1$ agents. The utility profile is:

Agents	a	b	c
$1, \dots, n$	1	0	0
$n + 1, \dots, 2n + 1$	0	$1 - \epsilon$	ϵ
Social Welfare	n	$(n + 1)(1 - \epsilon)$	$(n + 1)\epsilon$

The strategy profile in Nash equilibrium:

Agents	a	b	c
$1, \dots, n$	1	0	0
$n + 1$	0	$1 - \epsilon$	ϵ
$n + 2, \dots, 2n + 1$	0	0	1
Total	n	$1 - \epsilon$	$n + \epsilon$

The winner is alternative c . Observe that no agent can increase her utility by deviating. The PoA is $\frac{(n+1)(1-\epsilon)}{(n+1)\epsilon} = \frac{1-\epsilon}{\epsilon} \rightarrow \infty$ as $\epsilon \rightarrow 0$.

We now give an example for which there is no PNE. There are 2 alternatives $\{a, b\}$ and 2 agents. The mechanism is deterministic and ties are broken in favour of a . The utility profile is given below:

Agents	a	b
1	0.75	0.25
2	0.25	0.75

None of the three exhaustive cases below allow an equilibrium.

- If $u_1(a) < u_2(b)$ then b is the winner. For voter 1, this is not an equilibrium because she can increase her utility for a and make a win.
- If $1 > u_1(a) \geq u_2(b)$ then a is the winner. Now, for voter 2 this is not an equilibrium because she can increase her utility for b and make b win.

- If $u_1(a) = 1$ then a is the winner. Voter 2 cannot make b win, and as voters are truth-biased she will give her true input. Now, voter 1 is also truth-biased, and will give her true input, and a will remain the winner. This puts us in the second case above. ◀

We now show that the PoA even restricted to strong Nash equilibria is unbounded.

Proof of Theorem 21. We will give an example for which the price of anarchy $\rightarrow \infty$. There are 4 alternatives and $4n - 1$ agents. Ties are broken lexicographically. The utility profile is given below:

Agents \rightarrow	a	b	c	d
$1, \dots, 2n - 1$	1	0	0	0
$2n, \dots, 3n - 1$	0	$1 - \epsilon$	0	ϵ
$3n, \dots, 4n - 1$	0	0	$1 - \epsilon$	ϵ
Social Welfare	$2n - 1$	$n(1 - \epsilon)$	$n(1 - \epsilon)$	$2n\epsilon$

The strong pure Nash equilibrium strategy profile is:

Agents	a	b	c	d
$1, \dots, 2n - 1$	1	0	0	0
$2n, \dots, 3n - 1$	0	0	0	1
$3n, \dots, 4n - 2$	0	0	0	1
$4n - 1$	0	0	$1 - \epsilon$	ϵ
Total	$2n - 1$	0	$1 - \epsilon$	$2n - 1 + \epsilon$

The winner is alternative d . Observe that no agent can increase her utility by deviating. The PoA is $\frac{2n-1}{2n\epsilon} \rightarrow \infty$ as $\epsilon \rightarrow 0$. ◀

Iterative Voting

We give the proofs related to the PoA and convergence of iterative voting dynamics. We first restate the restrictions on dynamics from the main paper.

- (A) The utility of the deviating player must strictly increase after the deviation.
- (B) The deviating player can increase the reported utility for a single alternative, and this alternative must be chosen by the mechanism after the deviation.
- (C) The deviating player can decrease the reported utility for a single alternative, and this alternative must be chosen by the mechanism before the deviation.

We first show that just restriction (A) is insufficient to ensure bounded PoA.

► **Theorem 23.** *The price of anarchy of iterative voting with restriction (A) is unbounded.*

Proof. There are 5 alternatives $\{a, b, c, d, e\}$, and $2n + 4$ agents. Ties broken lexicographically. $\epsilon \rightarrow 0$ is much smaller than $1/n$. The utility profile is as given.

Agents	a	b	c	d	e
A: $1, \dots, n$	0.5	$0.5 - \epsilon$	ϵ	0	0
B: $n + 1, \dots, 2n$	$0.5 - \epsilon$	0.5	0	ϵ	0
$2n + 1, 2n + 2$	0	0	ϵ	0	$1 - \epsilon$
$2n + 3, 2n + 4$	0	0	0	ϵ	$1 - \epsilon$
Social Welfare	$n - n\epsilon$	$n - n\epsilon$	$(n + 2)\epsilon$	$(n + 2)\epsilon$	$4 - 4\epsilon$

27:20 On the Welfare of Cardinal Voting Mechanisms

Initially, alternative a is the winner. But an agent in B , say agent $n + 1$, can deviate and make b win. Say the agent deviates to: $0.5 + \epsilon$ to b , $0.5 - \epsilon$ to d , and 0 to others. Now, an agent in A , say agent 1, deviates and gives input: $0.5 + \epsilon$ to a , $0.5 - \epsilon$ to c , and 0 to others. Now, alternative a wins again. Repeating this for all agents in A and B , after n steps the strategy profile is given below. Currently, a is the winner.

Agents	a	b	c	d	e
$1, \dots, n$	$0.5 + \epsilon$	0	$0.5 - \epsilon$	0	0
$n + 1, \dots, 2n$	0	$0.5 + \epsilon$	0	$0.5 - \epsilon$	0
$2n + 1, 2n + 2$	0	0	ϵ	0	$1 - \epsilon$
$2n + 3, 2n + 4$	0	0	0	ϵ	$1 - \epsilon$
Total	$n(0.5 + \epsilon)$	$n(0.5 + \epsilon)$	$n(0.5 - \epsilon) + 2\epsilon$	$n(0.5 - \epsilon) + 2\epsilon$	$4 - 4\epsilon$

Now, the agent $2n + 3$ makes a move and sets d 's score to 1, then the agent $2n + 1$ makes a move and sets c 's score to 1. The same moves are then repeated by agents $2n + 4$ and $2n + 2$. This makes c the current winner. The strategy profile after these deviations is given below.

Agents	a	b	c	d	e
$1, \dots, n$	$0.5 + \epsilon$	0	$0.5 - \epsilon$	0	0
$n + 1, \dots, 2n$	0	$0.5 + \epsilon$	0	$0.5 - \epsilon$	0
$2n + 1, 2n + 2$	0	0	1	0	0
$2n + 3, 2n + 4$	0	0	0	1	0
Total	$n(0.5 + \epsilon)$	$n(0.5 + \epsilon)$	$n(0.5 - \epsilon) + 2$	$n(0.5 - \epsilon) + 2$	0

Alternatives a and b cannot win the election by a deviation by the agents in A or B , so the agents in A and B start competing for c and d to reach the final equilibrium strategy profile given below. Alternative c is the final winner.

Agents	a	b	c	d	e
$1, \dots, n$	0	0	1	0	0
$n + 1, \dots, 2n$	0	0	0	1	0
$2n + 1, 2n + 2$	0	0	1	0	0
$2n + 3, 2n + 4$	0	0	0	1	0
Total	0	0	$n + 2$	$n + 2$	0

No agent has a move that can increase her utility, and hence this is an equilibrium. The PoA is $\frac{n(1-\epsilon)}{(n+2)\epsilon} \rightarrow \infty$. Observe that this proof works even for best response iterative voting dynamics: The deviating agent plays a move that makes the most preferred alternative win the election, among the alternatives that can win the election after a move by the agent. ◀

With just restriction (A), as shown, we obtain unbounded PoA. Theorem 22 shows that with both (A) and (B) we get a near-optimal bound on the PoA. We now show, however, that the two restrictions (A) and (B) that ensured small PoA are not enough to ensure convergence to a Nash equilibrium.

► **Theorem 24.** *Iterative voting with restrictions (A) and (B) may not converge to a PNE.*

Proof. We will give a utility profile with 4 alternatives $A = \{a, b, c, d\}$ and 4 agents $N = \{1, 2, 3, 4\}$, and a sequence of steps taken by the agents that will create a cycle. For ease of writing, we normalize the utilities to sum up to 6 rather than 1. In the tables below the winner is denoted by *. The steps are:

1. Utility profile is

Alternatives →	a^*	b	c	d
Agent 1	1	1	0	4
Agent 2	2	2	0	2
Agent 3	0	2	3	1
Agent 4	4	0	2	0

2. Agent 1 plays. Utility increases from 1 to 4.

Alternatives →	a	b	c	d^*
Agent 1	0	0	0	6
Agent 2	2	2	0	2
Agent 3	0	2	3	1
Agent 4	4	0	2	0

3. Agent 4 plays. Utility increases from 0 to 2.

Alternatives →	a	b	c^*	d
Agent 1	0	0	0	6
Agent 2	2	2	0	2
Agent 3	0	2	3	1
Agent 4	0	0	6	0

4. Agent 1 plays. Utility increases from 0 to 1.

Alternatives →	a	b^*	c	d
Agent 1	0	5	0	1
Agent 2	2	2	0	2
Agent 3	0	2	3	1
Agent 4	0	0	6	0

5. Agent 3 plays. Utility increases from 2 to 3.

Alternatives →	a	b	c^*	d
Agent 1	0	5	0	1
Agent 2	2	2	0	2
Agent 3	0	0	5	1
Agent 4	0	0	6	0

6. Agent 4 plays. Utility increases from 2 to 4.

Alternatives →	a^*	b	c	d
Agent 1	0	5	0	1
Agent 2	2	2	0	2
Agent 3	0	0	5	1
Agent 4	6	0	0	0

7. Agent 3 plays. Utility increases from 0 to 2.

Alternatives →	a	b^*	c	d
Agent 1	0	5	0	1
Agent 2	2	2	0	2
Agent 3	0	2	3	1
Agent 4	6	0	0	0

27:22 On the Welfare of Cardinal Voting Mechanisms

8. Agent 1 plays. Utility increases from 1 to 4.

Alternatives \rightarrow	a	b	c	d^*
Agent 1	0	0	0	6
Agent 2	2	2	0	2
Agent 3	0	2	3	1
Agent 4	6	0	0	0

9. Agent 4 plays. Utility increases from 0 to 2.

Alternatives \rightarrow	a	b	c^*	d
Agent 1	0	0	0	6
Agent 2	2	2	0	2
Agent 3	0	2	3	1
Agent 4	0	0	6	0

Observe that the strategy profile in 3 and 9 are same, giving us a cycle. This proof also works for best-response dynamics. ◀