Incentive Compatibility, Condorcet, and Borda

BRAZENLY PRELIMINARY

Dean Jens

April 16, 2022

Abstract

Two important voting systems are the Condorcet method(s) and the Borda count. In this paper it is shown that these are two endpoints of a continuum in which the Condorcet method corresponds to higher levels of information (and associated incentive compatibility constraints) and the Borda count corresponds to lower levels of information.

Contents

1 Introduction **2**

2 Background **5**

2.1 Some Notation 5

2.2 Maximal Lotteries and Condorcet 6

2.3 Borda count 7

3 Results **7**
3.1 Party game 7
3.2 Voter epistemics 10
3.3 A couple of additional notes 12

1 Introduction

When a group of voters is to collectively decide between two alternatives, it is straightforward to have each voter indicate which alternative the voter prefers, and for the group to choose the alternative that is preferred by more voters. When there are more than two alternatives, the situation is less straightforward. Gibbard (1973); Satterthwaite (1975); Gibbard (1978) demonstrate that there are *no* election rules of general applicability for more than two candidates that do not violate at least one of a set of seemingly desirable criteria: that more than one of the voters can, at least under some circumstances, affect the outcome of the vote; that at least three of the candidates can, for some preferences, be chosen; and that the voting system is not subject to "strategic voting". This last criterion is a demanding one, requiring that a voter's best response to that voter's own preferences and the other voters' actions depend only on the former and not on the latter; as a practical matter, in any realistic voting system, a voter best optimizes by conditioning his or her vote on that voter's belief about what other voters are doing.

There are various approaches to softening the bite of this result. Dasgupta and Maskin (2008) restrict the domain of preference profiles. The largest domain on which it is then possible to construct a nontrivial non-strategic voting rule is the set of preference profiles for which there is a Condorcet winner, on which set Condorcet's method is such a rule. Any other rule fails unless it is restricted to a proper subset of that set of preference profiles. Saari (1990) takes a different approach, quantifying the "manipulability" of a voting rule and seeking the rule that minimizes it; depending somewhat on the number of candidates, his results favor the Borda count.[1] With greater attention to preferences with indifference between candidates, Brams and Fishburn (1978) show that approval voting is strategy-proof if all voters have preferences that divide the candidates into two subsets, within either of which the voter is indifferent; further, "approval voting is the only single-ballot nonranked voting system that has this property."

Condorcet's method[2] and the Borda count are both ranked voting systems — a vote consists of a list of the candidates in an order, typically interpreted as a preference order. (Dasgupta and Maskin, 2019, theorem 1) note that a strategy-proof, decisive voting system must be equivalent to a ranked voting system. "Strategy-proof" in this context is a Nash-type condition in a deterministic setting: a voter effectively knows which

[1]He also notes that just about any voting system minimizes some appropriately crafted measure of "manipulability", though he argues for his measure.

[2]Or "methods", depending on how a winner is chosen when there is no Condorcet winner.

candidate will win as a function of that voter's own vote, and accordingly the only thing that matters about that voter's preferences is which of any set of candidates is that voter's favorite. The present paper will consider players in games with less information about each others' play; the voters will see themselves as having a choice of distributions over candidates, so that preferences over lotteries can matter. (Dasgupta and Maskin, 2019, theorem 6) shows that, for three candidates on a domain of preferences that is sufficiently rich (a weak requirement), Condorcet's method and the Borda count are the only strategy-proof voting rules satisfying the Pareto principle, Anonymity (symmetry among voters), Neutrality (symmetry among candidates), and Decisiveness.

The context of the present paper will present a continuum of voting systems, with strategic information varying along the continuum. At one end, where it is common knowledge how voters are voting, is the "maximal lottery", a non-deterministic variant of the notion of "Condorcet winner" that will be presented in §2.2. It is a distribution over candidates; if there is a Condorcet winner, it puts full weight on the Condorcet winner. At the other end of the continuum, where there is very little information available about what voters are doing, we will find the Borda count.

One simple reduced-form manner of treating agents who lack full strategic information is the quantal response equilibrium of McKelvey and Palfrey (1995); Hofbauer and Sandholm (2002) note that such an agent with choice set A can be modeled as choosing a distribution $\sigma \in \Delta A$ over available actions so as to maximize the expected payoff from σ plus some concave function on ΔA that might be regarded as a cognitive or information cost associated with assiduously avoiding those actions that are not strictly optimal; indeed, if that concave function is the "entropy" of Shannon (1948), the result is the logistic quantal response equilbrium that is of special interest in the literature. This sort of model will be used in this paper in two different ways. One is an abstract, two-player game, presented in §3.1, in which two parties seek to maximize the number of votes by which they would win in a two-candidate election — as noted at the beginning, two-candidate settings are much simpler than multicandidate settings. The other context in which this is used is more inchoate at this point, and looks at how individual voters are incentivized to behave in the equilibria of such a game. This leads ultimately to the promised continuum of voting systems, trading expressiveness for incentive compatibility in as economical a way as possible.

2 Background

2.1 Some Notation

Let C be a finite set of **candidates** and I a set of voters. S_C is the set of permutations on C, $S_C = \{(a_1, \cdots, a_{|C|})|a_i \in C \forall i, i \neq j \Rightarrow a_i \neq a_j\}$. If an agent $i \in I$ has a **payoff** $u_i : C \to \mathbb{R}$, then an **ordinal preference** $s \in S_C$ is **consistent with** u_i if $i \leq j \Rightarrow u_i(s_j) \geq u(s_k)$. Given an ordinal preference

$s \in S_C$ we will define $r_s : C \to \{1, \cdots, |C|\}$ such that $r_s(s_i) = i$ for i in the codomain of r_s. If each voter has an ordinal preference, i.e. we have $s_i \in S_C \forall i \in I$, then, for any $p, q \in C$, let $m(p,q) = |\{i \in I | r_{s_i}(p) > r_{s_j}(p)\}| - |\{i \in I | r_{s_i}(p) < r_{s_j}(p)\}|$ denote the "margin" of p over q; note that, were such preferences interpreted as the true preferences of voters voting in a two-candidate election between p and q, this is the margin by which p would defeat q. We extend m bilinearly to pairs of distributions on C, $m : \Delta C \times \Delta C \to \mathbb{R}$, with no risk of confusion with the initial definition of m; indeed, there will be no care taken to distinguish $c \in C$ from the distribution in ΔC that places full weight on c.

2.2 Maximal Lotteries and Condorcet

Suppose each voter has an ordinal preference, so that "margins" are defined on pairs of distributions of candidates. Define a relation on the set of distributions of candidates as follows: for $p, q \in \Delta C$, let $p \succ q$ if the margin $m(p,q) > 0$. This relation always has a maximal element, called the **maximal lottery**.(Fishburn (1984)) Note that if there is a Condorcet winner[3] $c^* \in C$, then $m(c^*, c') > 0$ for all $c^* \neq c' \in C$, and thus the unique maximal lottery will place full weight on the Condorcet winner.

Brandl, Brandt, and Seedig (2016) show that the maximal lottery rule is the unique implementable rule that has an in-

[3]This very preliminary draft of this paper assumes the reader knows what a Condorcet winner is.

tuitive consistency under subdivision of sets of voters and under subdivision of sets of candidates; for example, if the electorate consists of the disjoint union of electorates that separately choose the same winner, then the electorate as a whole will choose that winner.

2.3 Borda count

The **Borda count** is a voting rule in which each voter i has an ordinal preference s_i, and the winning candidate $c \in C$ is chosen to maximizes $\sum_{c' \in C} m(c, c')$. It is more often defined to choose $c \in C$ to maximize some positive affine transform of $-\sum_{i \in I} r_{s_i}(c)$, but the definition in terms of margins will prove more convenient here. The definitions are equivalent; the set of maximizers of $\sum m(c, c')$ is equal to that of $-\sum r_s(c)$, even when they are not singleton, i.e. when there is a tie.

3 Results

3.1 Party game

Suppose there are two political parties, the Row Party and the Column Party. Neither party has an ideology or an agenda; each party just wants to win, and preferably to win big. In fact, each party seeks to maximize not its chances of winning, but the expected value of its margin of victory, and it will endorse a candidate to optimize its goal. If both parties endorse the same candidate, they both get a payoff of zero, but if they endorse different candidates, those two candidates are subject to a vote, and the party that chose the winning candidate gets a positive payoff equal to the vote margin, while the other party gets the negative of that.

The players of this game are the two parties; the pure strategy spaces are the set of candidates; and their payoffs are given by the margin m. The **party game** is the mixed extension, and is a two-player, zero-sum game.

With complete information, a two-player, zero-sum game is not epistemically demanding; players in such a game do not have to worry too much about what the other player is planning, or about what each player thinks the other player thinks. In this game, however, while the outsider knows the payoffs, the players do not. The parties may have various information about the voters' preferences, but there is a cost to acquiring information, and so the parties may occasionally be seen to choose a suboptimal candidate. In particular, let $S : \Delta(C) \to \mathbb{R}$ represent the negative of the cost of information; it may be the Shannon entropy, but we require only that it have the following properties:

- S is strictly concave
- ∇S diverges on the boundaries $\partial\Delta(C)$; in particular, approaching the relative interior of the face where $\sigma_c = 0$ for $\sigma \in \Delta(C)$, $c \in C$, we require $\partial S/\partial\sigma_c \to \infty$.

We will admit an additional parameter, $\lambda \in [0, \infty]$. In the quantal response equilibrium with $\lambda \in [0, \infty)$, the row player is observed to play the distribution $\sigma \in \Delta(C)$ that maximizes

$$u(\sigma) = \sum_c \sigma'_c m(\sigma, \sigma') + \lambda S(\sigma) \tag{1}$$

given the opposing player's choice of $\sigma' \in \Delta(C)$ and the column player plays the distribution $\sigma' \in \Delta(C)$ to maximize

$$u(\sigma') = \sum_c \sigma_c m(\sigma', \sigma) + \lambda S(\sigma')$$

taking σ as given. We allow $\lambda = \infty$, understanding $[0, \infty]$ topologically as the one-point compactification of $[0, \infty)$; there $u = S$ for both players.

Proposition 1. *If $\lambda = 0$, then in the equilibrium of the party game, both players are playing the maximal lottery.*

Proposition 2. *As $\lambda \to 0$, then in the equilibrium of the party game, both players are playing a distribution in which the winner of the Borda count is given the most weight.*

Proof. Sketch of proof: (Rockafellar, 1970, theorem 23.5) implies that ∇u is strictly monotonic in the important sense, McKelvey and Palfrey (1995); Hofbauer and Sandholm (2002) tell us that the equilibrium is unique and continuous in a neighborhood of ∞. Some first-order conditions and some algebra finish the proof. □

Now, consider the situation in which one party plays the quantal response equilibrium strategy for a given λ and the other party best-responds to it, i.e. chooses $p \in \Delta(C)$ to maximize $m(p, p^*)$ where p^* is the quantal response equilibrium strategy. In the case of $\lambda = 0$, this will result in the choice of a distribution over candidates that have support in the maximal lottery; in particular, if there is a Condorcet winner, that winner will necessarily be chosen. In the case of $\lambda = \infty$, p^* gives equal weight to all candidates, and the maximizing p will choose the winner of the Borda count.

3.2 Voter epistemics

The theorems of Gibbard (1973); Satterthwaite (1975); Gibbard (1978) concern voting mechanisms that are incentive compatible when voters are perfectly informed of each other's votes. Especially in settings with many voters, this is intuitively unreasonable; we might expect an information structure more like that of Myerson (1998), where voters know a probability distribution from which voter preferences are drawn, introducing some Poisson-distributed uncertainty into voters' knowledge of realized preferences. With ever less information, the incentive-compatibility constraints become, in an important sense, ever weaker; players can not respond to other players' actions if they do not know them.

Incentive compatibility (in the Gibbard-Satterthwaite context) and independence of irrelevant alternatives (in the Arrow context) are closely linked;[4] in particular, if a voter knows certain candidates are "irrelevant" but can still affect the outcome, the voter can feign preferences about that candidate to influence which of two other candidates is elected. When voters

[4] See e.g. Reny (2001).

are very well informed, an incentive-compatible voting system accordingly implies a set of candidates who are relevant, and whether a candidate is relevant becomes dependent only on the voters' preferences among the relevant candidates — a recursive problem that, in the case of a Condorcet winner, is solved only by the Condorcet winner, who does well against every set. When voters are more poorly informed, more candidates may seem somewhat relevant. The quantal response equilibria from the party game indicate the relevant candidates.

If voters have a great deal of uncertainty, and in particular if they put equal probability on any rank order by the other players, then the Borda count, which allows for a great deal of preference expressiveness, can become incentive compatible; the voter will expect, conditional on the vote, that the candidate listed highest has the best chance of winning, the candidate listed second second, and so on, and the voter thus optimizes by casting a ballot that gives the voter's ordinal preference. If the voter knows many candidates are irrelevant, the voter does well to place preferred relevant candidates at the top of the ballot, dispreferred relevant candidates at the bottom, and irrelevant candidates in the middle. The preference expressiveness, in practice, goes away, as the voter expresses false preferences.

While (as cited above) the Borda count is reported not to be "too" manipulable, informed agents are very unlikely to vote a ranked list that echoes the true preference order, and the Nash equilibria tend to be very sensitive to player beliefs. In a typical Nash equilibrium of a Borda count, or even a Myerson

(1998)-style Poission equilibrium, the relevant candidates are very close to tied, and voters who have somewhat better information about which candidates are actually within one voter's ability to determine the winner could make good use of it.

The p^* that obtains with a finite λ puts weight on those "relevant" candidates; the effect of a voter's vote on p^* will not depend on the position of irrelevant candidates in the voter's vote, because the relevant terms in (1) go to zero. Thus a voting system can be constructed to calibrate λ to some estimate of how informed voters are likely to be, find the quantal response equilibrium, and either use the equilibrium strategy to conduct a lottery or, perhaps more plausibly, elect the candidate who performs best against that lottery; this voting system can retain incentive compatibility while also, as much as possible, retaining expressiveness.

3.3 A couple of additional notes

Initial research into the effect of information on voter behavior in these contexts indicates that if voter preferences are entirely unconstrained, they can be crafted in such a way as to make voter incentives complicated, and indeed perhaps this should not be surprising given the behavior of the Borda count. Indeed, if the Borda count, in the limit as $\lambda \to \infty$, is supposed to represent a welfare measure that argues for this form of "preference expressiveness", it would have to be that the ordinal preference indicates all that we need to know about voter payoff, which suggests a restriction on voter preferences. Work remains to be done there.

The relationship of this work to (Dasgupta and Maskin, 2019, theorem 6) is somewhat weakened by that result's restriction to three candidates. Note that "neutrality" in the context of this paper's model is essentially a symmetry condition on S, and "anonymity" is imposed in the definition, though in principle weights could be incorporated into the definition of "margin".

In certain qualitative ways, intermediate values of λ can give rise to a setting that resembles approval voting. In particular, the optimal response to beliefs in approval voting consists of voting for those candidates who are preferred to the voter's perceived lottery (that is, the voter's payoff for each possible candidate, averaged with a weight proportional to the voter's subjective probability that that candidate will win) and against the other candidates; in the setting of this paper, the optimal ranked vote will rank the "approved" candidates above the "disapproved" candidates; as long as p^* does a good job of approximating voters' beliefs, the outcome is similar, especially if there is enough information that few of the candidates are "relevant".

References

Brams, S. J., and P. C. Fishburn (1978): "Approval Voting," *American Political Science Review*, 72(3), 831–847.

Brandl, F., F. Brandt, and H. G. Seedig (2016): "Consistent Probabilistic Social Choice," *Econometrica*, 84, 1839–1880.

Dasgupta, P., and E. Maskin (2008): "On The Robustness of Majority Rule," *Journal of the European Economic Association*, 6(5), 949–973.

——— (2019): "Elections and Strategic Voting: Condorcet and Borda," Discussion paper.

Fishburn, P. C. (1984): "Probabilistic Social Choice Based on Simple Voting Comparisons," *Review of Economic Studies*, 51(4), 683–692.

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

——— (1978): "Straightforwardness of Game Forms with Lotteries as Outcomes," *Econometrica*, 46(3), 595–614.

Hofbauer, J., and W. H. Sandholm (2002): "On the Global Convergence of Stochastic Fictitious Play," *Econometrica*, 70(6), 2265–2294.

McKelvey, R. D., and T. R. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10(1), 6–38.

Myerson, R. B. (1998): "Population uncertainty and Poisson games," *International Journal of Game Theory*, 27(3), 375–392.

Reny, P. J. (2001): "Arrow's theorem and the Gibbard-Satterthwaite theorem: a unified approach," *Economics Letters*, 70(1), 99–105.

Rockafellar, R. T. (1970): *Convex Analysis*. Princeton University Press.

Saari, D. G. (1990): "Susceptibility to Manipulation," *Public Choice*, 64(1), 21–41.

Satterthwaite, M. A. (1975): "Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions," *Journal of Economic Theory*, 10, 187–217.

Shannon, C. E. (1948): "A Mathematical Theory of Communication," *Bell System Technical Journal*, 27(3), 379–423.