Computational Social Choice: Autumn 2010

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Plan for Today

The broad aim for today is to show how we can *characterise* voting procedures via their properties.

We will give examples for three approaches:

- *Axiomatic method:* to characterise a (family of) voting procedure(s) as the only one satisfying certain axioms
- Distance-based approach: to characterise voting procedures in terms of a notion of *consensus* (elections where the outcome is clear) and a notion of *distance* (from such a consensus election)
- Voting as *truth-tracking:* to characterise a voting procedure as computing the most likely "correct" winner, given n distorted copies of an objectively "correct" ranking (the ballots)

Approach 1: Axiomatic Method

Two Alternatives

When there are only *two alternatives*, then all the voting procedures we have seen coincide, and *intuitively* they do the “right” thing.

Can we make this intuition precise?

- Yes, using the axiomatic method.

Anonymity and Neutrality

Recall how we have defined the properties of anonymity and neutrality of a voting procedure F:

- F is *anonymous* if $F(b_1, \ldots, b_n) = F(b_{\pi(1)}, \ldots, b_{\pi(n)})$ for any ballot profile $(b_1, \ldots, b_n)$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$.
- F is *neutral* if $F(\pi(\underline{b})) = \pi(F(\underline{b}))$ for any ballot profile $\underline{b}$ and any permutation $\pi : \mathcal{X} \to \mathcal{X}$ (with π extended to ballot profiles and sets of alternatives in the natural manner).

Positive Responsiveness

A voting procedure satisfies the property of *positive responsiveness* if, whenever some voter raises a (possibly tied) winner x in her ballot, then x will become the *unique* winner. Formally:

> F satisfies positive responsiveness if $x \in F(\underline{b})$ implies $\{x\} = F(\underline{b'})$ for any alternative x and any two distinct profiles $\underline{b}$ and $\underline{b'}$ with $\underline{b}(x \succ y) \subseteq \underline{b'}(x \succ y)$ and $\underline{b}(y \succ z) = \underline{b'}(y \succ z)$ for all alternative y and z different from x.

Remark: This is slightly stronger than *weak monotonicity*, which would only require $x \in F(\underline{b'})$. (Note that last week we had defined weak monotonicity for *resolute* voting procedures only.)

Notation: $\underline{b}(x \succ y)$ is the set of voters ranking x above y in profile $\underline{b}$.

May's Theorem

Now we can fully characterise the plurality rule:

Theorem 1 (May, 1952) *A voting procedure for two alternatives satisfies anonymity, neutrality, and positive responsiveness if and only if it is the plurality rule.*

Remark: In these slides we assume that there are no indifferences in ballots, but May's Theorem also works (with an appropriate definition of positive responsiveness) when ballots are weak orders.

K.O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 20(4):680–684, 1952.

Proof Sketch

Clearly, plurality does satisfy all three properties. ✓

Now for the other direction:

For simplicity, assume the number of voters is odd (no ties).
Plurality-style ballots are fully expressive for two alternatives.
Anonymity and neutrality $\rightsquigarrow$ only number of votes matters.

Denote as A the set of voters voting for alternative a and as B those voting for b. Distinguish two cases:

- Whenever $|A| = |B| + 1$ then only a wins. Then, by PR, a wins whenever $|A| > |B|$ (that is, we have plurality). ✓
- There exist A, B with $|A| = |B| + 1$ but b wins. Now suppose one a-voter switches to b. By PR, now only b wins. But now $|B'| = |A'| + 1$, which is symmetric to the earlier situation, so by neutrality a should win $\rightsquigarrow$ contradiction. ✓

Characterisation Theorems

When there are more than two alternatives, then different voting procedures are really different. To choose one, we need to understand its properties: ideally, we get a *characterisation theorem*.

Maybe the best known result of this kind is Young's characterisation of the *positional scoring rules* (PSR) ...

- Every scoring vector $s = \langle s_1, \ldots, s_m \rangle$ with $s_1 \geqslant s_2 \geqslant \cdots \geqslant s_m$ and $s_1 > s_m$ defines a PSR: give s_i points to alternative x whenever someone ranks x at the ith position; the winners are the alternatives with the most points.
- A *generalised* PSR is like a PSR, but without the constraint that $s_1 \geqslant s_2 \geqslant \cdots \geqslant s_m$ and $s_1 > s_m$.

Reinforcement (aka. Consistency)

A voting procedure satisfies *reinforcement* if, whenever we split the electorate into two groups and some alternative would win in both groups, then it will also win for the full electorate.

For a full formalisation of this concept we need to be able to speak about a voting procedure F wrt. different electorates $\mathcal{N}$, $\mathcal{N}'$, ...

To accommodate this need, we temporarily switch to a framework where voting procedures are explicitly parametrised by the electorate:

$$F^{\mathcal{N}} : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \to 2^{\mathcal{X}} \setminus \{\emptyset\}$$

We can now formally state the reinforcement axiom:

> F satisfies reinforcement if $F^{\mathcal{N} \cup \mathcal{N}'}(\underline{b}) = F^{\mathcal{N}}(\underline{b}|_{\mathcal{N}}) \cap F^{\mathcal{N}'}(\underline{b}|_{\mathcal{N}'})$ for any disjoint electorates $\mathcal{N}$ and $\mathcal{N}'$ and any ballot profile $\underline{b}$ (over $\mathcal{N} \cup \mathcal{N}'$) such that $F^{\mathcal{N}}(\underline{b}|_{\mathcal{N}}) \cap F^{\mathcal{N}'}(\underline{b}|_{\mathcal{N}'}) \neq \emptyset$.

Notation: $\underline{b}|_{\mathcal{N}} = (i \mapsto b_i \mid i \in \mathcal{N})$ is the profile of ballots by voters in $\mathcal{N}$ as given in the (possibly larger) ballot profile $\underline{b}$.

Continuity

A voting procedure is *continuous* if, whenever electorate $\mathcal{N}$ elects a unique winner x, then for any other electorate $\mathcal{N}'$ there exists a number k s.t. $\mathcal{N}'$ together with k copies of $\mathcal{N}$ will also elect only x.

(This is a very weak requirement.)

Young's Theorem

We are now ready to state the theorem:

Theorem 2 (Young, 1975) *A voting procedure satisfies anonymity, neutrality, reinforcement, and continuity if and only if it is a generalised positional scoring rule.*

Proof: Omitted (and difficult).

But it is not hard to verify the right-to-left direction.

H.P. Young. Social Choice Scoring Functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.

Approach 2: Consensus and Distance

Dodgson Rule

In 1876, Charles Lutwidge Dodgson (aka. Lewis Carroll, the author of *Alice in Wonderland*) proposed the following voting procedure:

- The *score* of an alternative x is the minimal number of pairs of adjacent alternatives in a voter's ranking we need to *swap* for x to become a *Condorcet winner*.
- The alternative(s) with the *lowest score* win(s).

A natural justification for this procedure is this:

- For certain ballot profiles, there is a clear *consensus* who should win (here: consensus $=$ existence of a Condorcet winner).
- If we are not in such a consensus profile, we should consider the closest consensus profile, according to some notion of *distance* (here: distance $=$ number of swaps).

What about other notions of consensus and distance?

Characterisation via Consensus and Distance

A generic method to define (or to "rationalise") a voting procedure:

- Fix a class of *consensus profiles:* ballot profiles in which there is a clear (set of) winner(s). (And specify *who* wins.)
- Fix a metric to measure the *distance* between two ballot profiles.
- This induces a *voting procedure:* for a given ballot profile, find the closest consensus profile(s) and elect the corresponding winner(s).

Useful general references for this approach are the papers by Meskanen and Nurmi (2008) and by Elkind et al. (2010).

T. Meskanen and H. Nurmi. Closeness Counts in Social Choice. In M. Braham and F. Steffen (eds.), *Power, Freedom, and Voting*, Springer-Verlag, 2008.

E. Elkind, P. Faliszewski, and A. Slinko. Distance Rationalization of Voting Rules. Proc. COMSOC-2010.

Notions of Consensus

Four natural definitions for what constitutes a consensus profile $\underline{b}$:

- *Condorcet Winner:* $\underline{b}$ has a Condorcet winner x ($\leadsto$ x wins)
- *Majority Winner:* there exists an alternative x that is ranked first by an absolute majority of the voters ($\leadsto$ x wins)
- *Unanimous Winner:* there exists an alternative x that is ranked first by all voters ($\leadsto$ x wins)
- *Unanimous Ranking:* all voters report the same ballot/ranking ($\leadsto$ the top alternative in that unanimous ranking wins)

(Other definitions are possible.)

Ways of Measuring Distance

Two natural definitions of distance between ballot profiles $\underline{b}$ and $\underline{b}'$:

- *Swap distance:* minimal number of pairs of adjacent alternatives that need be swapped to get from $\underline{b}$ to $\underline{b}'$.

  Equivalently: distance between two ballots = number of pairs of alternatives with distinct relative ranking (aka. *Kendall tau distance*); sum over voters to get distance between two profiles

  $$\sum_{i \in \mathcal{N}} \#\{(x,y) \in \mathcal{X}^2 \mid \{i\} \cap \underline{b}(x \succ y) \neq \{i\} \cap \underline{b}'(x \succ y)\}$$

  (Strictly speaking, this will be *twice* the swap distance.)

- *Discrete distance:* distance between two ballots is 0 if they are the same and 1 otherwise; sum over voters to get profile distance

  $$\#\{i \in \mathcal{N} \mid b_i \neq b_i'\}$$

More Ways of Measuring Distance

Other definitions of distance between ballot profiles are possible:

- other ways of measuring distance between individual ballots
- other ways (than sum-taking) of aggregating distances over voters
- even arbitrary metrics defined on pairs of profiles directly

  However, Elkind et al. (2010) show that the latter is too general to be useful: essentially *any* procedure is distance-rationalisable under such a definition.

E. Elkind, P. Faliszewski, and A. Slinko. On the Role of Distances in Defining Voting Rules. Proc. AAMAS-2010.

Examples

Two voting procedures for which the standard definition is already formulated in terms of consensus and distance:

- Dodgson Rule = Condorcet Winner + Swap Distance
- Kemeny Rule = Unanimous Ranking + Swap Distance

How about other procedures? Borda? Plurality?

Writings of C.L. Dodgson. In I. McLean and A. Urken (eds.), *Classics of Social Choice*, University of Michigan Press, 1995.

J. Kemeny. Mathematics without Numbers. *Daedalus*, 88:571–591, 1959.

Characterisation of the Borda Rule

Recall that the Borda rule is the PSR with vector $\langle m{-}1, m{-}2, \ldots, 0\rangle$.

Proposition 1 (Farkas and Nitzan, 1979) *The Borda rule is characterised by the unanimous winner consensus criterion and the swap distance.*

Proof sketch: The swap distance between a given ballot that ranks x at position k and the closest ballot that ranks x at the top is $k{-}1$. Thus, if voter i ranks x at position k she gives her $-(k{-}1)$ points. This corresponds to the PSR with vector $\langle 0, -1, -2, \ldots, -(m{-}1)\rangle$, which is equivalent to the Borda rule. ✓

Remark: So Dodgson, Kemeny, and Borda are all characterisable via the same notion of distance!

D. Farkas and S. Nitzan. The Borda Rule and Pareto Stability: A Comment. *Econometrica*, 47(5):1305–1306, 1979.

Characterisation of the Plurality Rule

Recall the plurality rule is the PSR with scoring vector $\langle 1, 0, \ldots, 0 \rangle$.

Proposition 2 (Nitzan, 1981) *The plurality rule is characterised by the unanimous winner consensus criterion and the discrete distance.*

Proof: Immediate.

Remark 1: to be precise, Nitzan used a slightly different distance

Remark 2: also works with Majority Winner + discrete distance, but doesn't work with Condorcet Winner or Unanimous Ranking

S. Nitzan. Some Measures of Closeness to Unanimity and their Implications. *Theory and Decision*, 13(2):129–138, 1981.

Application: Complexity of Winner Determination

Characterising voting procedures in terms of consensus and distance can be useful for proving results for entire families of voting procedures.

One promising direction is the work of Elkind et al. (2010), who establish upper bounds on the complexity of the *winner determination problem* (WDP):

- The WDP for any voting procedure characterised in terms of one of the four consensus criteria we have seen and a "reasonable" notion of distance (see paper) is in Θ_2^p ("parallel access to NP").
- Dodgson and Kemeny are Θ_2^p-complete, so they are worst cases in this respect (Hemaspaandra et al., 1997, Hemaspaandra et al., 2005).

E. Elkind, P. Faliszewski, and A. Slinko. On the Role of Distances in Defining Voting Rules. Proc. AAMAS-2010.

E. Hemaspaandra, L.A. Hemaspaandra, and J. Rothe. Exact Analysis of Dodgson Elections. *Journal of the ACM*, 44(6):806–825, 1997.

E. Hemaspaandra, H. Spakowski, and J. Vogel. The Complexity of Kemeny Elections. *Theoretical Computer Science*, 349:382–391, 2005.

Approach 3: Voting as Truth-Tracking

Voting as Truth-Tracking

An alternative interpretation of "voting":

- There exists an objectively "correct" ranking of the alternatives.
- The voters want to identify the correct ranking (or winner), but cannot tell with certainty which ranking is correct. Their ballots reflect what they believe to be true.
- We want to estimate the most likely ranking (or winner), given the ballots we observe. Can we use a voting procedure to do this?

Example

Consider the following scenario:

- two alternatives: A and B
- either $A \succ B$ or $B \succ A$ (we don't know which)
- 20 voters/experts with probability 75% each of getting it right

Now suppose we observe that 12/20 voters say $A \succ B$.
What can we infer, given this observation (let's call it X)?

- Probability for this to happen given that $A \succ B$ is correct:

$$P(X|A \succ B) \quad = \quad \binom{20}{12} \cdot 0.75^{12} \cdot 0.25^{8}$$

- Probability for this to happen given that $B \succ A$ is correct:

$$P(X|B \succ A) \quad = \quad \binom{20}{8} \cdot 0.75^{8} \cdot 0.25^{12}$$

$P(X|A \succ B)/P(X|B \succ A) \ = \ 0.75^4/0.25^4 \ = \ 81.$
Thus, given X, A being better is 81 times as likely as B being better.

The Condorcet Jury Theorem

For the case of *two alternatives*, the *plurality rule* (aka. simple majority rule) is attractive also in terms of truth-tracking:

Theorem 3 (Condorcet, 1785) *Suppose a jury of n voters need to select the better of two alternatives and each voter independently makes the correct decision with probability $p > \frac{1}{2}$. Then the probability that the plurality rule returns the correct decision increases monotonically in n and approaches 1 as n goes to infinity.*

Proof sketch: By the law of large numbers, the proportion of voters making the correct choice approaches $p \cdot n > \frac{1}{2} \cdot n$. ✓

For modern expositions see Nitzan (2010) and Young (1995).

Writings of the Marquis de Condorcet. In I. McLean and A. Urken (eds.), *Classics of Social Choice*, University of Michigan Press, 1995.

S. Nitzan. *Collective Preference and Choice*. Cambridge University Press, 2010.

H.P. Young. Optimal Voting Rules. *J. Economic Perspectives*, 9(1):51–64, 1995.

Characterising Voting Procedures via Noise Models

For n alternatives, Young (1995) has shown that if the probability of any voter to rank any given pair correctly is $p > \frac{1}{2}$, then the rule selecting the most likely winner coincides with the *Kemeny rule*.

Conitzer and Sandholm (2005) ask a general question:

- For a given voting procedure F, can we design a "*noise model*" such that F is a *maximum likelihood estimator* for the winner?

H.P. Young. Optimal Voting Rules. *J. Economic Perspectives*, 9(1):51–64, 1995.

V. Conitzer and T. Sandholm. Common Voting Rules as Maximum Likelihood Estimators. Proc. UAI-2005.

The Borda Rule as a Maximum Likelihood Estimator

It *is* possible for the Borda rule:

Proposition 3 (Conitzer and Sandholm, 2005) *If each voter independently ranks the true winner at position k with probability $\frac{2^{m-k}}{2^m-1}$, then the maximum likelihood estimator is the Borda rule.*

Proof: Let $r_i(x)$ be the position at which voter i ranks alternative x.

Probability to observe the actual ballot profile if x is the true winner:

$$\frac{\prod_{i \in \mathcal{N}} 2^{m-r_i(x)}}{(2^m-1)^n} = \frac{2^{\sum_{i \in \mathcal{N}} m-r_i(x)}}{(2^m-1)^n} = \frac{2^{\text{BordaScore}(x)}}{(2^m-1)^n}$$

Hence, x has maximal probability of being the true winner *iff* x has a maximal Borda score. ✓

V. Conitzer and T. Sandholm. Common Voting Rules as Maximum Likelihood Estimators. Proc. UAI-2005.

Summary

We have seen three approaches to characterising a voting procedure:

- as the only procedure satisfying certain *axioms*;
- as computing the closest *consensus* profile (wrt. some *distance*) with a clear winner and returning that winner; and
- as *estimating the most likely "true" winner*, given the information provided by the voters and some assumptions about how likely they are to estimate certain aspects of the "true" ranking correctly.

All three approaches are (potentially) useful

- to better understand particular voting procedures;
- to explain why there are so many "natural" voting procedures; and
- to help prove general results about families of voting procedures.

What next?

So far we have thought of voting procedures as functions that map profiles of *ballots* to sets of winners.

We did not (have to) speak about the *preferences* of voters.

Next we will connect these two levels and discuss *strategic* behaviour:

- Under what circumstances can we expect voters to vote *truthfully* (preference = ballot)?
- We will see that it is, in some technical sense, impossible to guarantee that no voter has an incentive to *manipulate* an election (by strategically choosing to provide an insincere ballot).
- Later we will discuss various ways of *circumventing* strategic manipulation (all of which can only be partial solutions).