Game Theory and Algorithms*
Lecture 16: Elections on a Line and in Practice

May 5, 2011

Summary: We continue our discussion of voting systems. One short example has to do with the power of "agenda-setters." We move on to define the Schulze method, which is becoming widely used in practice. Finally, we talk about ameliorating Arrow's theorem by moving in a different direction: what if we restrict the possible preference lists that voters can submit?

1 Agenda Control

Most businesses, societies and organizations follow some form of pairwise voting (using "Robert's Rules of Order.") The following example shows how this pairwise comparison can lead to very bad outcomes. Often, the chair who sets the agenda does not get a vote; nonetheless they may sometimes have more power than anyone else.

Consider a set of voters with the following preferences:

- 1/3 of voters have preferences $P > Q > A > Z$;
- 1/3 of voters have preferences $Q > A > Z > P$;
- 1/3 of voters have preferences $A > Z > P > Q$;
- (for descriptive purposes, imagine that Z is the favourite candidate of the agenda-setter).

Notice that absolutely every voter agrees Z is worse than A. However, it is still possible that Z can be declared the winner:

1. First, have a pairwise vote to see if A or Q is preferred. As 2/3 of the voters prefer Q to A, we will eliminate A.

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

2. Next, have a pairwise vote to see if P or Q is preferred. As 2/3 of the voters prefer P to Q, we will eliminate Q.

3. Finally, this leaves P and Z. In a pairwise vote to see if P or Z is preferred, Z wins with 2/3 of the votes.

Exercise. Draw a directed graph whose vertex set is the candidates; whenever A is preferred to B by a majority of voters, draw a directed edge from A to B. What conditions in this graph are necessary and sufficient for candidate C to be electable by iterated pairwise elimination votes?

2 The Schulze Method

The Schuzle Method is a particular voting method defined by giving strengths to different "paths" of comparisons. It was initially conceived in the late 1990s, and has grown quickly in popularity, starting in the mid-2000s. Notable adopters include the Linux projects Debian (2003) and Gentoo (2005), as well as various aspects of Wikipedia — the French language version in 2005, and the Wikimedia foundation in 2008. Unlike point-based systems, the Schulze method satisfies the Condorcet criterion, but (according to Wikipedia) it fails "participation" (whereas point-based systems satisfy participation). Schulze has published the method — see [2].

2.1 Definition

First, we compare every pair of candidates pairwise. For the sake of simplicity, assume as usual that we restrict candidates to strict linear orderings, although we can generalize the method to weak linear orders. Next, for every pair of candidates A, B, define $d[A, B]$ to be the number of voters that prefer A to B in pairwise comparisons. (So $d[A, B] + d[B, A] = n$.)

We now define the concept of a *strongest path*. Define a directed graph on the set of all candidates; and for each pair A, B draw a directed edge from A to B with label $d[A, B]$. We think of this label as the *strength* of the edge. Next, we extend the concept of strengths to paths in the following way:

the *strength* of a path equals the minimum strength of any edges it contains.

(This might bring to mind the expression, "a chain is only as strong as its weakest link.")

For each A, B, let $p[A, B]$ be the strength of the strongest path from A to B.

Exercise. Show that if $p[A, B] > p[B, A]$ and $p[B, C] > p[C, B]$, then $p[A, C] > p[C, A]$.

The output of the Schulze method is the ranking defined by

Output the order where A beats B if $p[A, B] > p[B, A]$.

As a result of the exercise above, it follows that this order satisfies transitivity. In words: the Schulze method ranks A above B if there is a "justification path" of pairwise beatings from A to B, such that the strength of voter approval on this path exceeds that of all justification paths from B to A.

Applying the same description to the overall result,

The Schulze method (ignoring ties) picks a winner W so that for every other candidate X, for some fraction $0 \leq f\% \leq 1$ of the population (depending on X), there is a pairwise chain of defeats from W to X supported by $f\%$ of the population, but no such chain from X to W.

Open-ended question: Is the "pairwise chain of defeats" a reasonable notion? For example, does anything bad happen similar to the critiques of backwards induction/centipede game?

Example 1. Take the following voters: 14 for $A > B > C > D$, 10 for $B > D > C > A$, 6 for $C > D > A > B$, 3 for $D > B > C > A$. The directed graph is shown below (except that for readability, we have left the labels off of some minority grey edges with weight less than $n/2 = 33/2$, which are not important for our calculations).

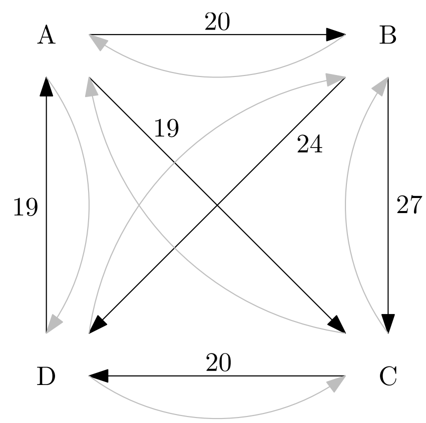

Let's compare A and B. The strongest path from A to B is the direct path consisting of just a single edge, and $p[A, B] = 20$. On the other hand, $p[B, A] = 19$, using the path $B \to D \to A$ for example. So $p[A, B] > p[B, A]$ and the mechanism would output A as beating B.

One other instructive calculation is $p[A, C]$. The strongest path is the indirect one via B, giving $p[A, C] = 20$. We get $p[C, A] = 19$ so C beats A.

Continuing these calculations, the overall output of the Schulze method is $A > B > C > D$.

Here are a couple of important good and bad properties of the Schulze method.

Proposition 2. *There is a polynomial-time algorithm for the Schulze method.*

The only part needing explanation is the computation of the strongest paths. This turns out to be already a pretty well-studied problem in graph theory, *widest paths*. The typical approach is that a *shortest path* algorithm can be easily modified to compute widest paths.

When $p[A, B] = p[B, A]$, the output of the Schulze method ranks A and B as tied.

Exercise. Assume that the labels $d[\cdot, \cdot]$ are distinct; no label appears more than once. (One imagines this is likely to occur when the number of voters is sufficiently large compared to the number of candidates.) Show that the Schulze method produces no ties: $p[A, B] \neq p[B, A]$ for all A, B.

Exercise. (a) Give an example with three candidates where the output of the Schulze method gives $A > B = C = A$ (so the Schulze method gives a quasi-transitive order, but not a weak linear order). (b) Give an example with four candidates where the output of the Schulze method gives $A > B = C > D = A$.

Even if one uses the Schulze method, there are a few parameters to be tweaked. First, one usually allows voters to submit weak partial orders; then we can define $d[A, B]$ in several ways (number of strictly supporting votes, number of weakly supporting votes, number of strictly supporting votes plus half of indifferent votes, ratio of supporting votes to opposed votes, etc). Second, when the output of Schulze method gives a tie, some further decision method must be specified (e.g., iterated deletion of all non-winners, flipping a coin; in the preliminary version of [2] Schulze suggests breaking ties according to a random voter, and iterating for any remaining ties).

Open-ended question: Is there a compelling reason to pick one method of defining d over the the other, or one method of breaking ties over the other?

3 Black's Condition

Given Arrow's theorem about the impossibility of certain voting systems, one alternative to relaxing pairwise independence is to look at scenarios where the choices of the voters are restricted.[1]

Definition 3. Given a *linearly ordered set of candidates* $\{1, 2, \ldots, k\}$, a voter's preference order O is *single-peaked* if for some i with $1 \leq i \leq k$,

$$1 <_O 2 <_O \cdots <_O i - 1 <_O i >_O i + 1 >_O \cdots >_O k - 1 >_O k.$$

Notice that there are no direct restrictions on the pairwise preferences between voters to the "left" and the "right" of the peak. For example, the "single-peaked" description $1 < 2 > 3 > 4$ applies to all of the following strict linear orders:

$$2 > 1 > 3 > 4, 2 > 3 > 1 > 4, 2 > 3 > 4 > 1.$$

[1]This motivation is not historically accurate: Black's theorem is from 1948, three year before Arrow's.

Theorem 4 (Black's single-peaked theorem)**.** *If all voters have single-peaked preferences, then there are no Condorcet cycles (so there is a Condorcet winner).*

Since Condorcet cycles are impossible, pairwise comparisons give a pairwise independent, unanimity-respecting, non-dictatorial mechanism. We will only prove the theorem in the case that the number of voters is odd. In the case that the number of voters is even, things are more complex (pairwise comparisons yield only a quasi-linear order, but still it can have no Condorcet cycle).

Proof. Since the number of voters is odd, any pairwise comparisons cannot result in a tie. Thus, if we model the pairwise winners by a directed graph, with a directed edge from A to B whenever A beats B pairwise, then every pair of nodes is linked by a directed edge in exactly one direction. This is usually called a *tournament* in graph theory.

Lemma 5. *A tournament has no cycles if and only if it has no cycles of length 3.*

Proof. Consider the *shortest* cycle in a tournament. We claim if any exists, it has length 3. This will complete the proof.

Let the shortest cycle be $v_1 \to v_2 \to \cdots v_t \to v_1$ and for the sake of contradiction assume $t > 3$. In what direction does the directed edge between v_1 and v_3 point? In either case (see Figure 1) a cycle of length smaller than t arises, giving the desired contradiction. □

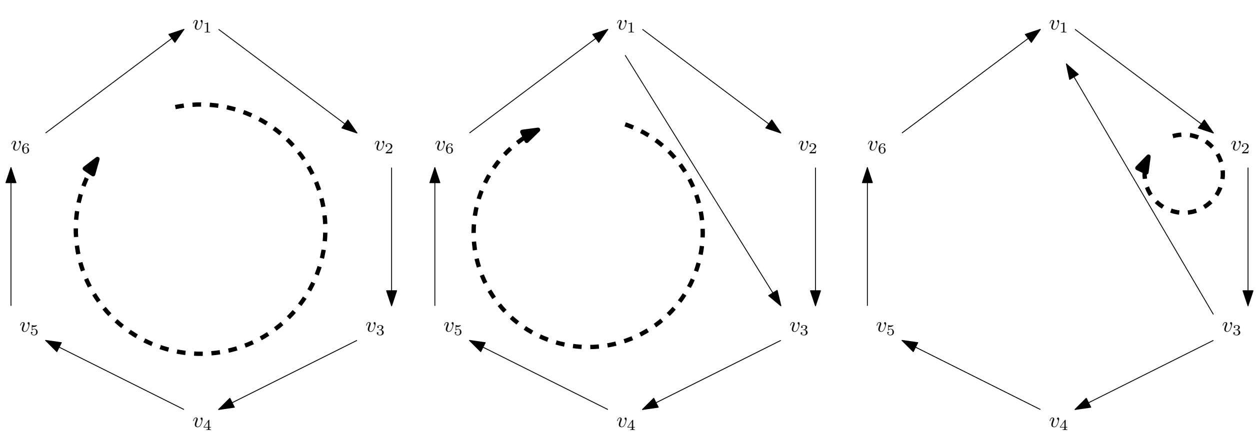

Figure 1: Left: a cycle of length t. Center: if there is an arc from $v_1 \to v_3$ we get another, smaller cycle. Right: if there is an arc from $v_3 \to v_1$ we get another, smaller cycle.

Lemma 6. *Single-peaked voters cannot create a Condorcet cycle of length 3.*

Proof. Suppose that there is a Condorcet cycle among 3 candidates. Without loss of generality (up to renaming and switching left-right) the cycle is $A \to B \to C \to A$ with A left, B middle, C right. We focus on these three candidates and use the following easy-to-check fact:

when we restrict single-peaked preferences to a subset of candidates, they are still single-peaked.

In the pairwise comparison between A and B, A wins. Notice this implies a majority of winners have their "peak" at A, since any other type of voter prefers B to A. But these same voters also prefer A to C, contradicting the edge $C \to A$. $\square$

These two lemmas, combined, give Black's theorem. $\square$

In fact, Black originally gave his theorem in a more constructive way. The following exercise helps you find Black's main idea, in a more general setting.

Exercise. Consider the case that the candidates are nodes on an undirected **tree** (rather than a line). A voter i is said to have *weakly single-peaked preferences* if for some node r_i, and all nodes y, and all nodes x between r_i and y on the tree, voter i's preferences amongst these pairs satisfy $r_i \geq_i x \geq_i y$. See Figure 2 for an example. Prove that there is a node C

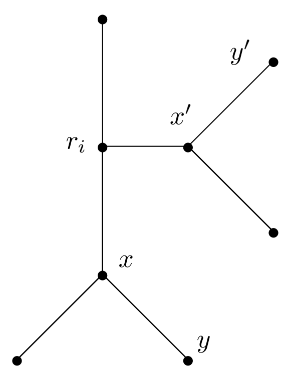

Figure 2: If a voter has weakly single-peaked preferences with peak r_i, and the tree of candidates is as shown, then voter i's preferences must include $r_i \geq x \geq y$ and $r_i \geq x' \geq y'$.

such that for all x, at least half of voters prefer $C \geq_i x$ (a sort of weak Condorcet winner). Hints: the location of C can be determined as a function of $r_1, r_2, \ldots, r_n$; it may be helpful to first solve the special case that the tree is a line.

Exercise (A false generalization)**.** If we allow weakly single-peaked preferences on a line, show Condorcet cycles can occur. Specifically, for three candidates on a line A, B, C, give weakly single-peaked preferences for some voters which give rise to a Condorcet cycle. (In this setting, directed edge $X \to Y$ means more voters prefer $X > Y$ pairwise than $Y > X$.)

Exercise (An algorithmic aside: Bartholdi & Trick)**.** If we are given the linear ordering of candidates, it is easy to check whether a given set of voter preferences are all single-peaked. However, what if we are not given the ordering of candidates? We could approach this by just testing all $k!$ possible orderings of candidates to see if any give rise to single-peaked preferences, but is there a polynomial-time algorithm? Show that the answer is "yes" by using a known polynomial-time subroutine [1] for the following problem:

The **Consecutive-Ones** Problem

Input: A matrix of 0s and 1s

Output: Does there exist a permutation of the columns, so that in every row, all of the 1s are consecutive?

For trees, the above exercise also admits a positive solution, using the *acyclic hypergraph* recognition algorithm of [3] in place of the consecutive-ones tester.

References

[1] G. S. Lueker and K. S. Booth. A linear time algorithm for deciding interval graph isomorphism. *J. ACM*, 26:183–195, April 1979.

[2] M. Schulze. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social Choice and Welfare*, 36:267–303, 2011. Preliminary version in *Voting matters*, 17:9–19, 2003.

[3] R. E. Tarjan and M. Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on Computing*, 13(3):566–579, 1984.