CS599: Structure and Dynamics of Networked Information (Spring 2005)
02/28/2005: Arrow's Theorem
Scribes: Shishir Bharathi, Fang Bian

1 Arrow's Social Choice Axiomatization

Last class, we followed the approach of Arrow [1], and derived a set of axioms that any "sensible" voting scheme should satisfy. These axioms were:

Monotonicity : If candidate A is ranked ahead of candidate B, and some voters previously preferring B over A change their ranking to place A ahead of B, then A still ends up ahead of B.

Non-triviality : For each pair A,B of candidates, there is some set of voter preferences that will cause the final ranking to have A precede B. In other words, the order of no pair is pre-determined.

Independence of Irrelevant Alternatives (IIA) : Changes in the preferences between some candidates do not affect the relative ordering of other candidates.

Non-dictatorship : the result of the voting should not just follow one individual's choice, while disregarding the others.

From monotonicity and non-triviality, we also derived the property of *unanimity*: if all voters favor candidate A over B, then A should be ranked ahead of B in the result.

2 Arrow's Theorem

Unluckily, with those four properties, we are asking for too much.

Theorem 1 (Arrow's Theorem (1951) [1]) *If there are at least three candidates, then there is no social choice function f satisfying all of the above axioms.*

Proof. To prove Arrow's theorem, we show that any social choice function satisfying monotonicity, non-triviality, and IIA is in fact a dictatorship function. We will do this by first proving the existence of a single voter who can decide the order between two candidates; then, we prove that this voter is in fact a dictator.

First, we define sets that decide the outcome between two candidates. We call a set J of voters *(A,B)-decisive* if the fact that all of J ranks A ahead of B is enough to guarantee that the A will be ranked ahead of B in the output. Formally, we write $A \prec_J B$ to denote that $A \prec_i B$ for all $i \in J$. We then say that J is (A,B)-decisive iff $A \prec_J B$ implies $A \prec B$. We call a set J of voters *decisive* iff J is (A,B)-decisive for some pair (A,B).

One useful characterization of decisiveness can be derived from the monotonicity property: J is (A,B)-decisive if and only if $A \prec_J B$ and $B \prec_{\bar{J}} A$ imply that $A \prec B$. We will use this characterization later. Notice also that by the unanimity property, the set V of all voters is always (A,B)-decisive for each pair (A,B). Hence, there must be some smallest decisive set J^*. We will prove that J^* is actually a singleton.

Consider an arbitrary voter $i \in J^*$ from this set, and define $J' = J^* - \{i\}$. Assume that the voter i has preference order CAB, each voter in J' has order ABC, and each voter in $\bar{J}^*$ has order BCA, for three candidates A, B, C. (Other candidates can be ignored by the IIA property). Then in the final order, A must precede B, because J^* is *(A,B)-decisive*, and all of J^* ranks A ahead of B. In addition, we know that C must precede A. The reason is that only J' would prefer A to precede C, so if it actually did, then J' would be (A,C)-decisive, contradicting the assumption that J^* is a smallest decisive set. Now, by transitivity, we conclude that the final output must be CAB, so C precedes B. But i is the only voter preferring C to precede B, so by our characterization above, $\{i\}$ must be (C,B)-decisive. Because $\{i\}$ is thus decisive, it cannot be smaller than J^*, so we know that $J^* = \{i\}$, and $\{i\}$ is also (A,B)-decisive.

Next, we show that $\{i\}$ is also (A, D)-decisive for all $D \neq A$ and (C, D)-decisive for all $D \neq C$. We consider the scenario where voter i has the order ABD, and everyone else has order BDA. Then, by unanimity, $B \prec D$, and because $\{i\}$ is (A, B)-decisive, we have $A \prec B$. By transitivity, the ordering is ABD. In particular, this means that A precedes D, which only voter i prefers. Thus, $\{i\}$ must be (A,D)-decisive. Replacing A by C in the previous proof shows that $\{i\}$ is also (C,D)-decisive. Also, by substituting C resp. A for D, we further conclude that $\{i\}$ must be (C,A)-decisive and (A,C)-decisive.

Next, we show that $\{i\}$ is also (D,A)-decisive for all $D \neq A$, as well as (D,C)-decisive for all $D \neq C$. Here, we assume that voter i prefers the order DCA, and everyone else prefers ADC. By unanimity, we have $D \prec C$ in the outcome, and because $\{i\}$ is (C,A)-decisive, we have $C \prec A$. So the final order will be DCA. Again, i is the only voter preferring D over A, so $\{i\}$ must be (D,A)-decisive. Since $\{i\}$ is both (A,C)-decisive and (C,A)-decisive, we can simply switch the order A and C in the previous construction, and prove that $\{i\}$ is (D,C)-decisive as well.

As a final step, we show that $\{i\}$ is in fact (D,E)-decisive for all D, E, and hence a dictator. We assume that i votes DAE, and everyone else votes EAD. Because $\{i\}$ is (D,A)-decisive and (A,E)-decisive, the final ordering must be DAE. But by the same argument as before, this means that $\{i\}$ is (D,E)-decisive.

In summary, we have proved that voter i is a dictator, completing our proof of Arrow's Theorem. ■

This result is quite disappointing. It suggests that for a very reasonable definition of what democratic decision-making is, the process is impossible. Yet, we observe frequently in practice that voting does work (reasonably) well. So one direction to pursue further is to ask: what kind of restrictions do voter preferences in practice seem to satisfy? Our construction in the proof was based on some very carefully crafted scenarios — perhaps, those don't appear in practice.

This line of thought is pursued further by Diane Richards, most importantly in [2]. One of the results in that paper is that if the candidates are all located on a line (corresponding to how far "left" or "right" their political opinions are), and each voter is a point on the line as well, sorting all candidates by increasing distance, then consensus voting is actually possible (i.e., Arrow's Axioms can be achieved).

3 Aggregation as Optimization

A different approach from the axiomatic one (which, in a sense, we just saw fail) would be to treat the determination of a consensus ordering as an optimization problem: to find a ranking that is as close as possible to the given set of rankings.

In order to phrase the problem this way, we first need to define a notion of distance between two sets of ranking. As rankings are really nothing but permutations (so long as they are complete rankings — more about partial rankings later), we can look at known distance measures on permutations. Two well-known such measures are Spearman's Footrule and Kendall's τ.

Definition 1 (Spearman's Footrule) *Let $\prec_1$, $\prec_2$ be two orderings, and $\prec_1 (A)$ the position of element A in the ordering $\prec_1$. Then, the Footrule distance is defined as*

$$F(\prec_1, \prec_2) \quad = \quad \sum_a | \prec_1 (a) - \prec_2 (a)|$$

Notice that this is identical to the L_1 distance between the vectors of positions.

Definition 2 (Kendall's τ) *Kendall's τ counts the number of inversions (or Bubble Sort swaps) between the two permutations. Writing*

$$K_{a,b}(\prec_1, \prec_2) \quad = \quad \begin{cases} 1 & \text{if } a \prec_1 b \text{ and } b \prec_2 a \\ 0 & \text{otherwise,} \end{cases}$$

we define

$$\tau(\prec_1, \prec_2) \quad = \quad \sum_{a,b} K_{a,b}(\prec_1, \prec_2).$$

The maximum value that can be attained by these metrics is $n(n-1)$ for Spearman's Footrule, and $\frac{n(n-1)}{2}$ for Kendall's τ. In both cases, the maximum is attained for two orderings that are the reverse of each other.

References

[1] K. Arrow. *Social Choice and Individual Values.* Wiley, 1951.

[2] D. Richards, B. McKay, and W. Richards. Collective choice and mutual knowledge structures. *Advances in Complex Systems*, 1:221–236, 1998.