Algorithms, Games, and Networks	February 12, 2013

Lecture 9

Lecturer: Ariel Procaccia — *Scribe: Jeremy Karp*

1 Overview

In the previous class, we discussed some basic ideas of social choice theory, including the definitions and properties of a number of voting rules. In today's class, we discuss manipulation in voting, with an emphasis on strategyproofness.

1.1 Review of the voting model

- Set of voters $N = \{1, \ldots, n\}$
- Set of alternatives A, $|A| = m$
- Each voter has a ranking over the alternatives
- $x \succ_i y$ means that voter i prefers x to y
- A **preference profile** is a collection of all voters' rankings
- A **voting rule** is a function from preference profiles to alternatives

Until today, we have assumed that voters are honest, they reveal their true preferences. Today, we will see that this may not always be an accurate assumption, as it is often in a voter's best interest to lie about his/her preferences.

1.2 A motivating example

1	2	3
b	b	a
a	a	b
c	c	c
d	d	d

Figure 1: In this profile, where all voters are truthful, b wins

1	2	3
b	b	a
a	a	c
c	c	d
d	d	b

Figure 2: By changing his vote so that it no longer is indicative of his true preferences, voter 3 is able to achieve a better outcome

Using the Borda count rule, it is possible for a voter to achieve a more favorable outcome by changing her vote away from her true preference. Consider the above figure for an example of this.

1.3 Strategyproofness

Intuitively, we call a voting rule strategyproof if it is not possible for a voter to benefit from lying about her preferences. Mathematically, we express this as follows:

$$\forall \prec,\ \forall i \in N,\ \forall \prec_i',\ f(\prec) \succeq_i f(\prec_i', \prec_{-i})$$

Consider the plurality voting rule. This rule is only strategyproof when $m \leq 2$, where m is the number of alternatives. We see this because when $m = 2$, plurality is strateyproof because if a voter's first choice is winning, they can't benefit from lying. Similarly, if the voter's favored alternative is losing, lying will not shift the vote in their favor. Now consider when $m \geq 3$. If a voter's third choice is in first place, voting for one's second choice alternative could break a tie in favor of this alternative, a preferable outcome compared the one's third choice winning the election.

A constant function is when the outcome of a vote is always the same, regardless of the voters' stated preferences. A dictatorship is when the winning alternative will be a single voter's top-ranked preference. Both of these voting rules are strategyproof. A constant function's outcome is unaffected by the voters' rankings, so not only is this rule strategyproof, the voters' preferences are entirely irrelevant. A dictatorship is also strategyproof: if a voter is not the dictator, one's preferences will not influence the vote at all, and if a voter is the dictator she receives the most utility from stating her true preferences.

2 Gibbard-Satterthwaite

A voting rule is **dictatorial** if there is a voter who always gets his preferred alternative, regardless of the other voters' stated preferences. A voting rule is **onto** if any alternative can win, for some set of stated preferences.

Theorem 1 (Gibbard-Sattherwaite): *If $m \geq 3$ then any voting rule that is strategyproof and onto is dictatorial*

We will use the following two lemmas to prove this theorem for $n = 2$:

Lemma 2 (Stong monotonicity): *f is a strategyproof rule, $\prec$ is a profile, and $f(\prec) = a$. Then $f(\prec') = a$ for all profiles $\prec'$ such that $\forall x \in A,\ i \in N :\ [a \succ_i x \Rightarrow a \succ_i' x]$*

Lemma 3 (Pareto optimality): *f is a strategyproof and onto rule, $\prec$ is a profile. If $a \succ_i b$ for all $i \in N$ then $f(\prec) \neq b$*

Proof: We only prove the theorem for the case of two voters. We begin our proof by defining two rankings (Figure 3), with the following properties:

$$\prec_1:\ a \succ b \succ x\ \forall x \in A \backslash \{a, b\}$$
$$\prec_2:\ b \succ a \succ x\ \forall x \in A \backslash \{a, b\}$$

$\prec_1$	$\prec_2$
a	b
b	a
c	d
d	c

Figure 3: Rankings $\prec_1$ and $\prec_2$

$\prec_1$	$\prec_2'$
a	b
b	d
c	c
d	a

Figure 4: Rankings $\prec_1$ and $\prec_2'$

By Pareto optimality, we see that $f(\prec) \notin A \setminus \{a, b\}$. Then, $f(\prec) \in \{a, b\}$. Without loss of generality, we say that $f(\prec) = a$. Now, we define a new ranking, $\prec_2'$ (Figure 4), with the following properties:

$$\prec_2':\ b \succ x \succ a\ \forall x \in A \backslash \{a, b\}$$

Similarly, we observe that $f(\prec') \notin A \setminus \{a, b\}$ by Pareto optimality. Furthermore, $f(\prec') \neq b$, otherwise voter 2 would be able to manipulate the outcome to his benefit, which would violate our assumption of strategyproofness. Then it is necessarily true that $f(\prec') = a$.

By strong monotonicity, any preference profile where voter 1 ranks a first would result in a winning the election. Then we say that voter 1 is "dictator for a".

Now, we need to show that voter 1 is the dictator for all alternatives, not just a. We define the set $A_i = \{x \in A : \text{ i is dictator for } x\},\ i \in \{1, 2\}$ and $A_3 = A \backslash A_1 \cup A_2$. We now observe that $|A_3| \leq 1$. If this were not the case, then we would have two alternatives $x, y \in A_3$. Then we could repeat the analysis in the previous paragraph, replacing a with x and b with y and we would find that one of the voters is dictator for one of these alternatives, contradicting that both x and y are elements of A_3.

It is also the case that $A_2 = \emptyset$, since either A_1 or A_2 must be empty. $A_3 = \emptyset$ as well because if $x \in A_3$ then we could repeat our original analysis with voter 1 preferring x and voter 2 preferring a and we would find that either $x \in A_1$ or $a \in A_2$. Either way, this is a contradiction. With $A_2 = A_3 = \emptyset$, we conclude that $A_1 = A$, proving that this voting rule is dictatorial. ■

3 Single Peaked Preferences

A municipality wants to choose a public location to build a library. In this model, the set of alternatives are the possible locations for the library. Each voter has a peak, their ideal location for the library. The closer the library is to a voter's peak, the happier he is.

Two possible voting rules are choosing the leftmost peak and choosing the midpoint of the peaks. The midpoint of the peaks is intuitively more appealing, but notice that this voting rule is not strategyproof, while choosing the leftmost peak is strategyproof. If we use the midpoint rule, assuming that a voter knows the location of the other voters' stated peaks, then the voter can alter his peak to shift the midpoint closer to his true peak. With the leftmost rule, a manipulator can only shift the selected peak further left. In any case, this will not improve the outcome for the manipulator.

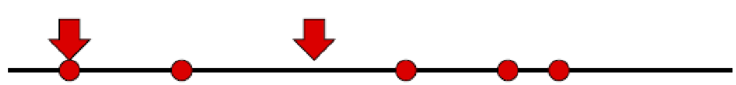

Figure 5: An example of several voters' peak preferences. The arrows indicate the location we would choose if we chose the leftmost and midpoint peaks, respectively.

The median is a compelling voting rule in this voting model. The median is a Condorcet winner, onto, and nondictatorial. The median is also strategyproof! If a voter's true peak is to the left of the median, the only way to alter the median is to report a location on the right of the median. This, however, hurts the voter so there is no incentive to do this. Similarly, if the voter's true peak is to the right of the median, the voter only worsens his outcome by lying about his preferred outcome. Of course, if a voter's preferred outcome is the median, then they cannot benefit from manipulating the outcome.

4 Complexity of Manipulation

With the Gibbard-Satterthwaite Theorem in mind, we know that all commonly used voting rules can be manipulated. However, even if it is possible in theory to manipulate a voting rule, it may still be computationally difficult to do. In this next section, we consider the possibility that some voting rules are easier to manipulate than others.

4.1 The R-Manipulation problem

The problem is the following: given votes of nonmanipulators and a preferred candidate p, can the manipulator cast vote that makes p uniquely win under R? One method that will successfully answer this question in many circumstances is a simple greedy algorithm.

This algorithm will rank p in first place. Then, while there are unranked alternatives the algorithm checks if there is an alternative that can be placed in the next spot without preventing p from winning. If there is such an alternative, we choose this to be our selection for the spot in question and move on to the next unfilled spot. If every unranked alternative prevents p from winning, then the algorithm returns "false," indicating that the greedy algorithm cannot find a ranking that successfully manipulates the outcome of the vote.

Theorem 4 (Bartholdi et al., SCW 89): *Fix $i \in N$ and the votes of other voters. Let R be a rule such that there exists a function $s(\prec_i, x)$ such that:*

- *For every $\prec_i$ chooses a candidate that maximizes $s(\prec_i, x)$*
- $\{y : y \prec_i x\} \subseteq \{y : y \prec_i' x\} \Rightarrow s(\prec_i, x) \leq s(\prec_i', x)$

Then the greedy algorithm always decides the R-Manipulation problem correctly.

4.2 Voting rules that are hard to manipulate

Several of the voting rules we are familar with are difficult to manipulate:

- Copeland with second order tie breaking
- STV
- Ranked Pairs