A Bit of Evolutionary Game Theory followed by Social Choice Theory and Voting

Jan van Eijck, CWI Amsterdam and Uil-OTS Utrecht
Rineke Verbrugge, Institute of AI, University of Groningen

ESSLLI 2009, Bordeaux, July 21, 2009

Overview

- A bit of evolutionary game theory
- A very brief history of voting theory
- Examples of voting procedures
- Arrow's Theorem
- A Proof of Arrow's Theorem, in Discourse Form
- What does it all mean?

A Bit of Evolutionary Game Theory: Hawks and Doves

- Evolutionary biologists Maynard Smith and Price [3] devised conflict model: members of a species fight over some resource.
- Winning the resource is worth 50 points.
- Players can use two strategies, “hawk” and “dove”.
- Hawk: fight over the resource.
- Dove: merely threaten and posture.
- If both play hawk, they will fight until one is injured (-100 points) and the other gets the resource.
- If both play dove, one will eventually get the resource of 50 points, but both get -10 points for wasted time.

Hawk and dove: the pay-offs

The pay-offs are computed as expected values:

- for hawk-hawk: $\frac{1}{2}(50) + \frac{1}{2}(-100) = -25$.
- for dove-dove: $\frac{1}{2}(50) - 10 = 15$.

	hawk	dove
hawk	(-25,-25)	(50,0)
dove	(0,50)	(15,15)

Note that this is again non zero-sum.

Hawk and Dove as Evolutionary Game

In populations with pure hawks and doves:

- If there are almost entirely doves, the doves expect on average 15 points per game. Then hawks are genetically advantaged: they expect an average of 50 points and their population would rise: the hawk minority would invade the dove population.
- A population of almost entirely hawks would be unstable: then the doves have an advantage, winning an average of 0 points against -25 points for hawks.
- A population of $\frac{7}{12}$ hawks and $\frac{5}{12}$ doves is stable, and the proportion stays fixed. In such a population, the expected pay-offs for both hawks and doves are $\frac{25}{4}$.

Questions

- Can the same equilibrium be attained if every individual plays the same mixed strategy of $\frac{7}{12}$ hawk and $\frac{5}{12}$ dove?

More generally:

- Our representations are quite abstract . . .
- How realistic are our scenarios. Do they apply to real life situations?

A Very Brief History of Voting Theory

A Very Brief History of Voting Theory

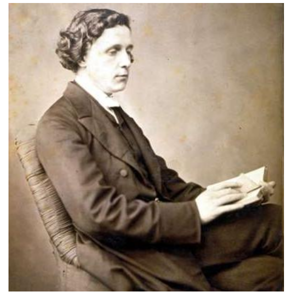

Jean-Charles de Borda 1733 – 1799

Marquis de Condorcet 1743 – 1794

Rev. Charles Dodgson 1832 – 1898

Kenneth Arrow b. 1921

Some Terminology

Terminology follows the excellent textbook Taylor [6]. Let there be a set A of alternatives and a set $1, \ldots, N$ of voters.

R is a binary relation on A if $R \subseteq A^2$.

R is transitive if xRy and yRz together imply xRz.

R is complete if at least one of xRy and yRx holds.

R is a weak ordering of A if R is a binary relation on A that is transitive and complete.

A ballot is a weak ordering of A.

A profile is a list $(R_1, \ldots, R_N)$ of ballots of A.

R_i represents the preferences of voter i. xR_iy expresses that x is at least as good as y according to voter i.

Ties

- Profiles represent sets of voter preferences, where ties are allowed.
- Let xPy abbreviate (xRy and not yRx). Then P expresses strict preference.
- Let xIy abbreviate (xRy and yRx). Then I expresses indifference or a tie.

Linear Ballots and Linear Profiles

R is anti-symmetric if xRy and yRx together imply that $x = y$.

R is a linear ordering of A if R is a binary relation on A that is transitive, complete and anti-symmetric.

A linear ballot is a linear ordering of A.

A linear profile is a list $(R_1, \ldots, R_N)$ of linear ballots of A.

Below we will use $(L_1, \ldots, L_N)$ to emphasise that the ballots are linear.

Linear profiles represent sets of voters preferences without ties.

Simple Majority Voting, or Plurality Voting

Each voter names his or her preferred candidate, and the candidate with the most votes is the winner. The winner of simple majority voting is called the plurality winner.

Borda Count, Borda Method, Borda Winner

A single-winner election method in which voters rank candidates in order of preference. The Borda count determines the winner of an election by giving each candidate a certain number of points corresponding to the position in which he or she is ranked by each voter. Once all votes have been counted the candidate with the most points is the winner (the Borda winner)

Let us assume that ties are not allowed. Then voter i has a linear ballot: a linear ordering of the list of candidates. Candidate x will get as i-Borda count the number of candidates below x in i's ranking. The total Borda count for x is the sum of the Borda counts for x for all the voters.

Example Outcome of Borda Voting Procedure

Let the profile be as follows:

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

Then a gets $2+0+0=2$, b gets $1+2+1=4$ and c gets $0+1+2=3$, so b is the Borda winner.

Note that for this profile there is no single majority for any alternative.

Condorcet's Voting Paradox

Assume there are three voters 1, 2, 3, and three alternatives a, b, c.

The profile is as follows:

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

Then:

- a beats b by two to one votes
- b beats c by two to one votes
- But: a is beaten by c by two to one votes.

This shows that majority voting is not transitive.

Condorcet Winner, Condorcet Method

The Condorcet winner of a profile is the choice that beats all alternatives in pairwise comparisons. The Condorcet voting paradox yields example profiles where there is no Condorcet winner.

The Condorcet method is the voting procedure where the Condorcet winner, if there is one, wins the election, and there is no winner of the election if no Condorcet winner exists.

Condorcet remarked that there are profiles where a Condorcet winner exists but does not coincide with the Borda winner.

[Exercise for today: find such a profile.]

Single Transferable Vote Method, or Hare's Method

The alternatives are voted for all at once. The one who gets the fewest votes is eliminated, and the process is repeated.

Also called the Hare method, after Thomas Hare who proposed it in 1861. One of its advocates was John Stuart Mill.

Social Choice Functions

- A social choice function is a function from linear profiles to alternatives (or: social choices).
- So voting systems are social choice functions.
- Arrow's result is a theorem about social choice functions satisfying certain conditions.
- A first natural condition is the Universal Domain Condition: social choice functions should be defined for any profile.

Pareto Efficient Social Choice Functions

- A social choice function is Pareto efficient if whenever social good x is at the top of every voter's preference list, then the function yields value x.
- This requirement is very reasonable. If the profile is unanymously in favour of x, x should be the outcome of the vote.

Monotonicity, or Independence of Irrelevant Alternatives

- A social choice function is monotonic if the following holds: if the function yields choice x for profile $L_1 \ldots L_N$, then the choice does not change if we adjust the preferences of all the voters, provided in each new preference $L_i{}'$ no social good y that was ranked below x in L_i is promoted to rank above x.
- This is also called independence of irrelevant alternatives.
- Suppose Rineke and Jan go to a restaurant with fish, meat and vegetarian on the menu. The restaurant has a constraint: one choice per table. They consult with each other, ..., and they decide to order the fish dish. The waiter then casually informs them that the meat dish is no longer available. Do they need another round of consultation then?

Dictatorial Social Choice Functions

- A social choice function is dictatorial if there is a voter i such that it holds for every input vector $(L_1 \dots L_N)$ that the social choice is x if and only if x is at the top of i's preference ranking L_i.
- Ideally social choice functions should be democratic in the sense that the weight of a vote does not depend on the voter: if f assigns x to the profile $(L_1, \dots, L_N)$, then f should also assign x to any permutation of that profile.
- Dictatorial social choice functions are not democratic ;-)
- Dictatorial social choice functions are bad ;-(

Arrow's Theorem, and Social Choice Theory

The field that is called Social Choice Theory was founded following the celebrated impossibility results that Kenneth Arrow proved in his thesis (later published in [1]).

> Any Pareto efficient and monotonic social choice function is dictatorial.

A Proof, in Discourse Form, taken from [2]

Logician: [...] I will give you a proof of the fact that any Pareto efficient and monotonic social choice function is dictatorial. From this the Gilbert-Satterthwaite result easily follows.

Computer Scientist: But first you have to explain to us what it means for a social choice function to be Pareto efficient and monotonic. Pareto efficiency, I can guess: a social choice function f is Pareto efficient if whenever social good x is at the top of every voter's preference list, then f yields value x.

Logician: That's right. Monotonicity is also straightforward. If social choice function f yields choice x for preference vector $L_1 \dots L_N$, then the choice does not change if we adjust the preferences of all the voters, provided in each new preference $L_i{}'$ no social good y that was ranked below x in L_i is promoted to rank above x.

Philosopher: So only lowering the position of x in the voter preferences might effect a change from x to a different choice. This is the Gibbard-Satterthwaite counterpart of independence of irrelevant alternatives, I suppose.

Logician: Indeed, it is. Now here is the theorem: if there are at least three social goods, and f is a social choice function that is Pareto efficient and monotonic, then f is dictatorial.

Philosopher: Fine. Let's go for the proof.

Logician: Suppose x, y are distinct social goods. Assume a voter profile with x top of the list and y bottom of the list in every voter's ranking. What should the outcome of f be?

Philosopher: Well, x, of course. This follows from the fact that f is Pareto efficient.

Logician: That's right. Remember that y was bottom of the list for every voter. Now suppose that I take the preference list of the first voter, and start moving y upward on the list. What will happen?

Computer Scientist: As long as y stays below x, nothing I suppose.

Logician: And if I move y above x?

Computer Scientist: Either nothing, or the value changes to y. This follows by monotonicity of f, doesn't it?

Logician: Correct. Now suppose I am going through the voter list, and for each voter move y from the bottom position to the top position. What will happen?

Philosopher: Then for some voter i, at the point where y gets raised past x, the choice will change from x to y. For suppose it does not. Then we end up with a preference list where y is above x in every voter's preference, while the choice is still x. This contradicts Pareto efficiency.

Logician: That's right. So we get the following two pictures. Let's call these Figure 1 and Figure 2. (Draws two pictures for them to look at.)

$$
\begin{array}{ccccccccc}
L_1 & \cdots & L_{i-1} & L_i & L_{i+1} & \cdots & L_N & & \\
y & \cdots & y & x & x & \cdots & x & & \\
x & \cdots & x & y & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & \mapsto & x \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & \cdot & y & \cdots & y & &
\end{array}
$$

This first picture shows the situation just before the value flips from x to y.

$$
\begin{array}{ccccccc}
L_1 & \cdots & L_{i-1} & L_i & L_{i+1} & \cdots & L_N \\
y & \cdots & y & y & x & \cdots & x \\
x & \cdots & x & x & \cdot & \cdots & \cdot \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot \\
\cdot & \cdots & \cdot & \cdot & y & \cdots & y
\end{array}
\quad \mapsto \quad y
$$

This second picture shows the situation just after the value has flipped from x to y.

Philosopher: Fair enough. And now I suppose further on in the proof the patterns in these pictures get manipulated a bit more?

Logician: That's exactly right. Let us study what would happen if in the first picture and the second picture we were to move x down to the bottom for all voters below i, and move x down to the second last position for all voters above i.

Computer Scientist: Nothing, I suppose.

Logician: That's right, the situations would be as pictured in the following figures. Let us call these Figures 3 and 4.

$$\begin{array}{ccccccccc}
L_1 & \cdots & L_{i-1} & L_i & L_{i+1} & \cdots & L_N & & \\
y & \cdots & y & x & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & y & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & \longmapsto & x \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & \cdot & x & \cdots & x & & \\
x & \cdots & x & \cdot & y & \cdots & y & &
\end{array}$$

This is Figure 3. It is the result of taking Figure 1 and moving x down to the bottom for voters below i and moving x to the second last position for voters above i.

$$
\begin{array}{ccccccccc}
L_1 & \cdots & L_{i-1} & L_i & L_{i+1} & \cdots & L_N & & \\
y & \cdots & y & y & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & x & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & \mapsto & y \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & \cdot & x & \cdots & x & & \\
x & \cdots & x & \cdot & y & \cdots & y & &
\end{array}
$$

This is Figure 4. It is the result of making similar changes to Figure 2.

Philosopher: For Figure 4, I can see why the value does not change. In Figure 2 the value was y, and it must remain y in Figure 4 by monotonicity.

Computer Scientist: OK, so Figure 4 has value y. But the Figures 3 and 4 differ only in the order of x, y in the ranking of i. It follows by monotonicity that the value in Figure 3 must be either y or x.

Philosopher: And it cannot be y, because then by monotonicity the value in Figure 1 would have to be y as well, and it is not. So the value in Figure 3 has to be x.

Logician: Just as I told you. Now suppose I take Figure 3 and move y down to the one but last position for all voters below i. This would not change the choice from x to a different value, would it?

Philosopher: I suppose it would not, by monotonicity again. The relative position of y with respect to x does not change.

Logician: So we get the following picture:

$$
\begin{array}{ccccccccc}
L_1 & \cdots & L_{i-1} & L_i & L_{i+1} & \cdots & L_N & & \\
\cdot & \cdots & \cdot & x & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & y & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & \mapsto & x \\
\cdot & \cdots & \cdot & \cdot & \cdot & \cdots & \cdot & & \\
y & \cdots & y & \cdot & x & \cdots & x & & \\
x & \cdots & x & \cdot & y & \cdots & y & &
\end{array}
$$

Call this Figure 5. Now consider a social good z different from x and y. By moving z through the preference orderings without letting z move past x we can obtain the following situation without changing the value of the choice function:

$$
\begin{array}{ccccccccc}
L_1 & \cdots & L_{i-1} & L_i & L_{i+1} & \cdots & L_N & & \\
\cdot & \cdots & \cdot & x & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & z & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & y & \cdot & \cdots & \cdot & \mapsto & x \\
z & \cdots & z & \cdot & z & \cdots & z & & \\
y & \cdots & y & \cdot & x & \cdots & x & & \\
x & \cdots & x & \cdot & y & \cdots & y & &
\end{array}
$$

Call this Figure 6.

Philosopher: I suppose monotonicity ensures that the value of the function does not change by the transition from 5 to 6?

Logician: That's correct. Now swap the rankings of x and y for all voters above i. By monotonicity, the choice value for the result must be either x or y.

Computer Scientist: But it cannot be y. For suppose it is, and consider the effect of moving z to the top in every preference. Since this would nowhere effect a swap with y, the value would have to remain y, by monotonicity. But then a profile with everywhere z on top would have value y, which contradicts Pareto efficiency.

Logician: So the value has to remain x, and we get the following picture:

$$\begin{array}{ccccccccc}
L_1 & \cdots & L_{i-1} & L_i & L_{i+1} & \cdots & L_N & & \\
\cdot & \cdots & \cdot & x & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & z & \cdot & \cdots & \cdot & & \\
\cdot & \cdots & \cdot & y & \cdot & \cdots & \cdot & \mapsto & x \\
z & \cdots & z & \cdot & z & \cdots & z & & \\
y & \cdots & y & \cdot & y & \cdots & y & & \\
x & \cdots & x & \cdot & x & \cdots & x & &
\end{array}$$

Now we are done, for observe that monotonicity ensures that making changes in the preferences of i while making sure that x remains on top will have no effect on the outcome. This means that the social choice will be x whenever x is at the top of i's ranking.

Philosopher: So i is a dictator for social good x. But since x was arbitrary, there must also be a dictator j for social good z as well. Clearly if i dictates whether x is on top, and j whether z is on top, then, to paraphrase Henk Wesseling, i and j have to be the same guy. Hence i must be a dictator for all alternatives.

How to Interpret Arrow's Results

- Illuminating view on Arrow's result: Don Saari, in [4, 5].
- Saari: the monotonicity constraint needs to be modified. Monotonicity, or independence of irrelevant alternatives, or binary independence, does not take the intensity of a preference of a voter for x over y into account.
- Saari proposed to replace I by what he calls the principle of intensity of binary independence. Let's call it II. This principle states that also the intensity of a voter's preference of one alternative over another should be taken into account.
- In particular, it matters not only that x is preferred over y, but also how many candidates there are between x and y.

Saari's Result

- Positional voting methods are methods that score candidates by allotting numbers of points to them to reflect their position in the preference ordering of a voter. The paradigm of this is the Borda count.
- Plurality voting can be viewed as a positional voting system: you assign 1 point for x for all ballots where x is at the top. Similarly for antiplurality voting.
- Saari has a theorem stating that the only positional social choice function satisfying U, P and II is the Borda count. All other positional methods fail.
- Arrow's theorem hinges on the fact that the principle I of independence of irrelevant alternatives allows one to hide the rationality of the voters.

Condorcet's Objection to Positional Methods

Condorcet's example (from [4]), with six groups of voters and three alternatives:

one	two	three	four	five	six
30	1	29	10	10	1
x	x	y	y	z	z
y	z	x	z	x	y
z	y	z	x	y	x

Any reasonable positional voting system will elect y, but this not the Condorcet winner!

Saari's Representation

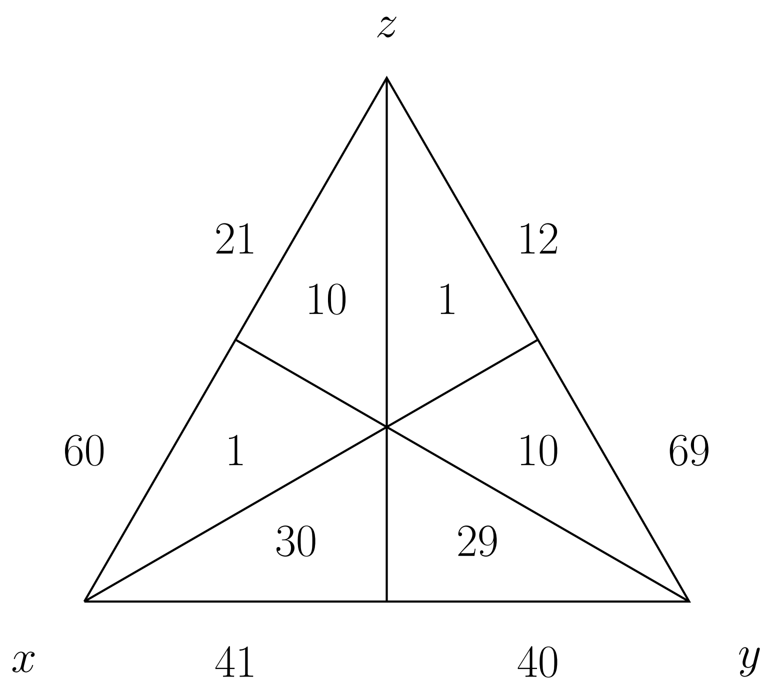

Simplifying Saari's Representation

- Look for groups of ballots that cancel out.
- Should the ballot pair $x < y < z$ and $z < y < x$ cancel out?
- How about the ballot triple $x < y < z$, $y < z < x$, $z < x < y$?
- There is one more such pattern, with a cycle in the other direction: $z < y < x$, $y < x < z$, $x < z < y$.

Ballot Patterns that Cancel Out

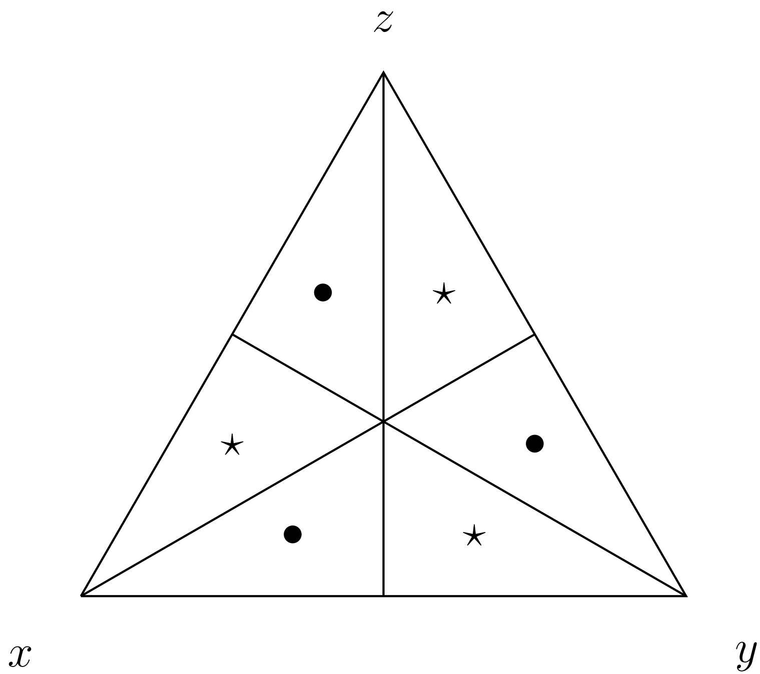

Result of Tie Simplification

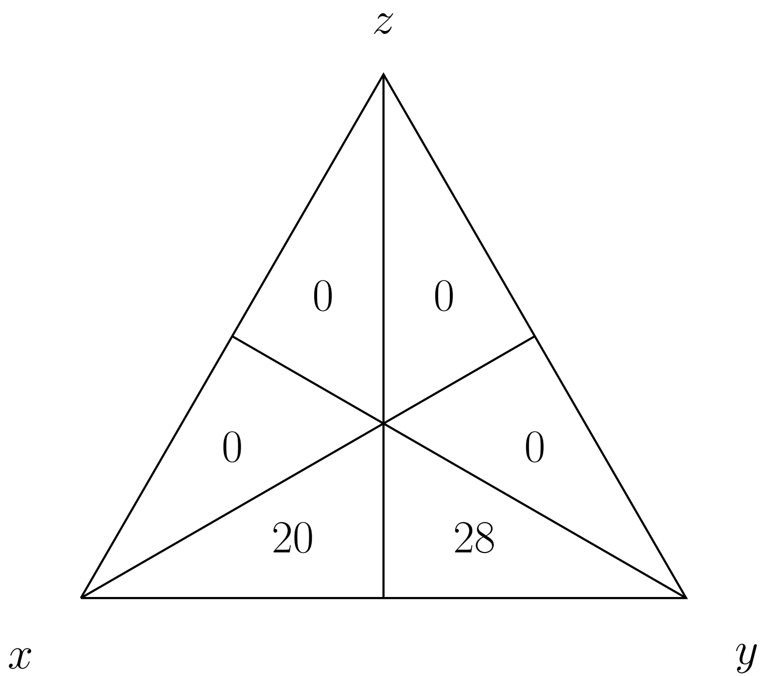

References

[1] K. Arrow. Social Choice and Individual Values. Wiley, New York, 1951, second edition: 1963.

[2] J. van Eijck and R. Verbrugge, editors. Discourses on Social Software, volume 5 of Texts in Logic and Games. Amsterdam University Press, Amsterdam, 2009.

[3] John Maynard Smith. Evolution and the Theory of Games. Cambridge University Press, 1973.

[4] Donald G. Saari. Chaotic Elections. American Mathematical Society, 2001.

[5] Donald G. Saari. Decisions and Elections: Explaining the Unexpected. Cambridge University Press, 2001.

[6] Alan D. Taylor. Social Choice and the Mathematics of Manipulation. Cambridge University Press, 2005.