Computational Social Choice and Manipulation in Approval Voting

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

Computational Social Choice

Social choice theory studies mechanisms for *collective decision making*, such as voting procedures or fair division protocols.

Computational social choice adds a computational perspective to this, and also explores the use of concepts from social choice in computing.

This Talk

First, three examples showing that computer science can offer a useful new perspective on problems in collective decision making:

- Computational Barriers against Manipulation in Voting
- Compact Representation of Preferences in Combinatorial Domains
- Computing Fair and Efficient Allocations of Goods to Agents

Then, one topic in more technical detail:

- Manipulation in Approval Voting

Problem: Vote Manipulation

Suppose the *plurality rule* (as in most real-world situations) is used to decide the outcome of an election: the candidate receiving the highest number of votes wins.

Assume the preferences of the people in, say, Florida are as follows:

$$
\begin{array}{ll}
49\%: & \text{Bush} \succ \text{Gore} \succ \text{Nader} \\
20\%: & \text{Gore} \succ \text{Nader} \succ \text{Bush} \\
20\%: & \text{Gore} \succ \text{Bush} \succ \text{Nader} \\
11\%: & \text{Nader} \succ \text{Gore} \succ \text{Bush}
\end{array}
$$

So even if nobody is cheating, Bush will win in a plurality contest.

Issue: In a pairwise competition, Gore would have defeated anyone.

Issue II: It would have been in the interest of the Nader supporters to *manipulate*, *i.e.* to misrepresent their preferences (and vote for Gore).

Approach: Make Manipulation Intractable

By the Gibbard-Satterthwaite Theorem, *any* voting rule for choosing between ≥ 3 candidates can be manipulated (unless it is dictatorial).

Idea: So it's always *possible* to manipulate, but maybe it's *difficult*! Tools from *complexity theory* can be used to make this idea precise.

- For the *plurality rule* this does *not* work: if I know all other ballots and want c to win, it is *easy* to compute my best strategy.
- But for *single transferable vote* it does work. Bartholdi and Orlin showed that manipulation of STV is *NP-complete*.

A. Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4):587–601, 1973.

M.A. Satterthwaite. Strategy-proofness and Arrow's Conditions. *Journal of Economic Theory*, 10:187–217, 1975.

J.J. Bartholdi III and J.B. Orlin. Single Transferable Vote Resists Strategic Voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

Problem: Huge Numbers of Alternatives

The alternatives often have a *combinatorial structure:* they are characterised by a tuple of variables ranging over a finite domain.

- allocate n indivisible goods to m agents: m^n alternatives
- elect a committee of size k, from n candidates: $\binom{n}{k}$ alternatives

Just *representing* and communicating the *preferences* of the agents can become a non-trivial problem (and that's not the only problem).

People in AI have long worked on *knowledge representation*, so there is a lot of expertise in this community . . .

Approach: Compact Representation of Preferences

We need languages that can represent preferences in a compact way. One type of language are so-called *weighted propositional formulas*. Utility is computed as the sum of the weights of the formulas satisfied.

Example: $\{(a,3),(b \vee c \vee d, 4),(b \wedge \neg c, 2)\}$ defines this utility function:

$$
\begin{array}{lll}
u(\emptyset) = 0 & u(ab) = 9 & u(abc) = 7 \\
u(a) = 3 & u(ac) = 7 & u(abd) = 9 \\
u(b) = 6 & u(ad) = 7 & u(acd) = 7 \\
u(c) = 4 & u(bc) = 4 & u(bcd) = 4 \\
u(d) = 4 & u(bd) = 6 & u(abcd) = 7 \\
 & u(cd) = 4 &
\end{array}
$$

Questions: *Expressivity*, relative *succinctness*, computational *complexity*?; how to use this for preference *aggregation*?; ...

J. Uckelman and U. Endriss. *Preference Representation with Weighted Goals: Expressivity, Succinctness, Complexity*. Proc. AiPref-2007.

Problem: Finding Socially Optimal Allocations

Scenario: Group of *agents* and set of indivisible *goods*. Agents have *preferences* over bundles of goods. What would be a good *allocation*?

Welfare economics and social choice theory give various definitions for what is *socially optimal* (e.g., utilitarianism *vs*. egalitarianism).

Problem: How do we find (compute) a socially optimal allocation?

Solution I: Computer scientists have developed powerful algorithms for computing socially optimal solutions in a *centralised* manner (integer programming, constraint satisfaction, heuristic search techniques).

Also interesting are *distributed* approaches . . .

Approach: Convergence in Distributed Negotiation

We have studied several variations of the following model:

- Preferences: agents have arbitrary quasi-linear utility functions.
- Agents will accept all deals that benefit them (and only those).
- No structural restrictions on deals (≥ 2 agents possible etc).
- Side payments are possible and agents have "enough" money.

Then a known result states that *any sequence of deals* will eventually converge to an allocation that has *maximal utilitarian social welfare*.

Questions: What structural restrictions on *deals* work for which types of *preferences*?; other notions of social optimality (*fairness*)?; how many deals before termination (*communication complexity*)?

U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of Artif. Intelligence Research*, 25:315–348, 2006.

Computational Social Choice

Computational social choice combines ideas from mathematical economics and computer science in new and fruitful ways.

For further information, have a look at our "*Short Introduction to Computational Social Choice*".

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. *A Short Introduction to Computational Social Choice*. Proc. SOFSEM-2007.

Outline of the Rest of the Talk

- Background: the Gibbard-Satterthwaite Theorem
- Background: Approval Voting
- Tie-Breaking and Preferences over Sets of Candidates
- Results: Manipulation in Approval Voting
- Conclusion

The Gibbard-Satterthwaite Theorem

Theorem 1 (Gibbard-Satterthwaite) *Every voting rule for three or more candidates must be either dictatorial or manipulable.*

Let C be a finite set of *candidates* and let $\mathcal{P}$ the set of all linear orders over C. A *voting rule* for n *voters* is a function $f : \mathcal{P}^n \rightarrow C$, selecting a *single winner* given the (reported) voter preferences.

A voting rule is *dictatorial* if the winner is always the top candidate of a particular voter (the dictator).

A voting rule is *manipulable* if there are situations where a (single) voter can force a preferred outcome by misreporting his preferences.

A. Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4):587–601, 1973.

M.A. Satterthwaite. Strategy-proofness and Arrow's Conditions. *Journal of Economic Theory*, 10:187–217, 1975.

The Gibbard-Satterthwaite Theorem

Theorem 1 (Gibbard-Satterthwaite) *Every voting rule for three or more candidates must be either dictatorial or manipulable.*

Despite its generality, the Gibbard-Satterthwaite Theorem may not apply in all cases (at least not immediately):

- The theorem presupposes that a ballot is a full preference ordering over all candidates. Plurality voting, for instance, does not satisfy this condition (although it's manipulable anyway).
- The theorem also presupposes that there is a *unique* way of casting a *sincere ballot* for any given preference ordering.

We will concentrate on the second "loophole". We can imagine various situations in which there may be more than one way of casting a sincere vote . . .

Approval Voting

In approval voting, a *ballot* is a subset of the set of candidates. These are the candidates the voter approves of. The candidate receiving the most approvals *wins* (we'll discuss tie-breaking later).

Approval voting has been used by several professional societies, such as the American Mathematical Society (AMS).

We assume each voter has a *preference ordering* $\preceq$ over candidates (which is antisymmetric, transitive and total).

A given voter's ballot is called *sincere* if all approved candidates are ranked above all disapproved candidates according to that voter's $\preceq$.

Example: If $A \succ B \succ C$, then $\{A\}$, $\{A, B\}$ and $\{A, B, C\}$ are all sincere ballots. The latter has the same effect as abstaining.

S.J. Brams and P.C. Fishburn. Approval Voting. *American Political Science Review*, 72(3):831–847, 1978.

Tie-breaking and Preferences over Sets of Candidates

We call the candidates with the most approvals the *pre-winners*.

If there are two or more pre-winners, we have to use a suitable *tie-breaking rule* to choose a winner.

Tie-breaking is outside the control of voters (in general). So when considering to manipulate, they have to do so in view of their preferences over sets of pre-winners.

- Given a voter's preferences $\preceq$ over *individual candidates*, what can we say about his preferences $\trianglelefteq$ over *sets of pre-winners*?
- Answer: *depends* (we'll consider several possible axiomatisations)

Axioms for Preferences over Sets of Candidates

Recall: Given a voter's preferences $\preceq$ over *individual candidates*, we want to characterise his preferences $\trianglelefteq$ over *sets of pre-winners*.

We can try to axiomatise the range of possible choices for $\trianglelefteq$ we wish to admit. One very reasonable option is this:

- $\trianglelefteq$ is reflexive and transitive.
- (DOM) $A \trianglelefteq B$ if $\#A = \#B$ and there exists a surjective mapping $f : A \to B$ such that $a \preceq f(a)$ for all $a \in A$.
- (ADD) $A \trianglelefteq B$ if $A \subset B$ and $a \preceq b$ for all $a \in A$ and all $b \in B \setminus A$.
- (REM) $A \trianglelefteq B$ if $B \subset A$ and $a \preceq b$ for all $a \in A \setminus B$ and all $b \in B$.

An Example for Successful Manipulation

Suppose all but one voter have voted. This final voter wants to manipulate. His preferences are: $4 \succ 3 \succ 2 \succ 1$.

Suppose 3 and 1 each got 10 votes so far (*pivotal* candidates); 4 and 2 each got 9 (*subpivotal* candidates). The final voter can

- force outcome 431 by voting [4];
- force outcome 3 by voting [43], [432], [3] or [32];
- force outcome 31 by voting [4321], [431], [321] or [31];
- force outcome 4321 by voting [42];
- force outcome 1 by voting [421], [41], [21] or [1]; or
- force outcome 321 by voting [2].

Outcomes 431, 3 and 4321 are *undominated* according to our axioms. If (and only if) the final voter prefers 4321 amongst these, he has an incentive to submit the insincere ballot [42].

The Case of Three Candidates

Recall that the Gibbard-Satterthwaite hits once we move from two to three candidates. The previous example shows that approval voting is certainly manipulable in the case of four candidates ...

Theorem 2 (Three candidates) *In approval voting with three candidates, suppose that all but one voter have cast their ballot. Then the final voter has no incentive to cast an insincere ballot.*

This is similar to a result by Brams and Fishburn (1978).

S.J. Brams and P.C. Fishburn. Approval Voting. *American Political Science Review*, 72(3):831–847, 1978.

Proof of Theorem 2

Check all possible cases. For each candidate, distinguish whether she is pivotal (P), subpivotal (S) or insignificant (I). At least one has to be pivotal, so there are $3^3 - 2^3 = 19$ possible situations.

```
?- table(3).
-----------------------------------------------------
|     | [100] [110] [111] | [001] [010] [011] [101] |
-----------------------------------------------------
|    [... 9 obvious cases of the form P__ omitted]   |
| SPP |  321   2     21   |  1     2     21    1    |
| SPS |  32    2     2    |  21    2     2     321  |
| SPI |  32    2     2    |  2     2     2     32   |
| SSP |  31    321   1    |  1     21    1     1    |
| SIP |  31    31    1    |  1     1     1     1    |
| IPP |  21    2     21   |  1     2     21    1    |
| IPS |  2     2     2    |  21    2     2     21   |
| IPI |  2     2     2    |  2     2     2     2    |
| ISP |  1     21    1    |  1     21    1     1    |
| IIP |  1     1     1    |  1     1     1     1    |   ✓
-----------------------------------------------------
```

The Case of Four Candidates

We know that manipulation is possible with four candidates (see earlier example). But *how many* problematic situations are there?

Answer: Just one!

Theorem 3 (Four candidates) *In approval voting with four candidates, suppose that all but one voter have cast their ballot. Then the final voter has no incentive to cast an insincere ballot, unless he strictly prefers* 4321 *over both* 431 *and* 3*.*

The proof has been derived automatically using a computer program that checks all possible scenarios.

The Case of Five Candidates

```
?- theorem(5).
Theorem: In approval voting with 5 candidates,
suppose that all but one voter have cast their ballot.
Then the final voter has no incentive to cast an
insincere ballot, unless his preferences over sets
of candidates satisfy one of the following 10 conditions:
-- 54321 strictly dominates all of 5431, 4.
-- 54321 strictly dominates all of 5421, 4.
-- 54321 strictly dominates all of 542, 4.
-- 5432 strictly dominates all of 542, 4.
-- 54321 strictly dominates all of 541, 4.
-- 5431 strictly dominates all of 541, 4.
-- 5421 strictly dominates all of 541, 4.
-- 54321 strictly dominates all of 531, 5431, 3.
-- 5321 strictly dominates all of 531, 3.
-- 4321 strictly dominates all of 431, 3.
```

Changing the Axioms

Next we assume that voters are *expected-utility maximisers*, but we want to drop any assumptions about the tie-breaking rule used.

Fully *general tie-breaking* can be axiomatised like this:

- (GEN) $A \trianglelefteq B$ if $a \preceq b$ for all $a \in A$ and all $b \in B$.

Question: How does this affect our manipulability results?

The Case of General Tie-Braking

By widening the range of conceivable orderings $\trianglelefteq$ over pre-winners, we have to give up our sincerity theorem for three candidates:

```
?- theorem(3). % based on (GEN) only
Theorem: In approval voting with 3 candidates,
suppose that all but one voter have cast their ballot.
Then the final voter has no incentive to cast an
insincere ballot, unless his preferences over sets
of candidates satisfy one of the following 2 conditions:
-- 321 strictly dominates all of 32.
-- 21 strictly dominates all of 31, 321.
```

For four candidates, there are already 19 such exceptions . . .

The Case of Uniform Tie-Breaking

If we are more specific about how ties are broken and how voters form preferences over sets of pre-winners we can obtain stronger results:

Theorem 4 (Expected-utility maximisers) *In approval voting with uniform tie-breaking, suppose that all but one voter have cast their ballot. Then, if the final voter is an expected-utility maximiser, he has no incentive to cast an insincere ballot.*

Remark: We need not make any assumptions regarding the actual utility functions used by the voters, other than that they are compatible with their preference orderings $\preceq$.

Discussion

What does this last result tell us about the merits of approval voting?

- Uniform tie-breaking in combination with the assumption that voters are expected-utility maximisers is arguably the most relevant scenario in practice. $\leadsto$ no need to be *insincere*
- But *manipulation* is still possible! Voters can strategise by choosing the most promising of their sincere ballots. We only show that there is no need for them to consider insincere ballots.
- The number of sincere ballots is *linear*; the number of insincere ballots is *exponential* in the number if candidates.

Conclusion

- Basic idea: The presence of *multiple sincere ballots* may allow us to circumvent the Gibbard-Satterthwaite Theorem in the sense that some sincere ballot may always be optimal.
- Results: For *approval voting*, it turns out that this is indeed the case for several interesting scenarios:
  - If all voters are *optimistic* or *pessimistic* (*not shown*).
  - If there are at most *three* candidates and one of these hold:
    - Voter preferences are governed by our "*reasonable axioms*".
    - A *rational chair* is breaking the ties (*not shown*).
  - If *uniform tie-breaking* is used and the voters are *expected-utility maximisers*.

U. Endriss. Vote Manipulation in the Presence of Multiple Sincere Ballots. *Proc. 11th Conference on Theoretical Aspects of Rationality and Knowledge*, 2007.