Vote Manipulation in the Presence of Multiple Sincere Ballots

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

Talk Outline

- Background: the Gibbard-Satterthwaite Theorem . . . and why there is hope that it may not apply in *all* cases
- Background: Approval Voting (with multiple sincere ballots)
- Tie-Breaking and Preferences over Sets of Candidates
- Results: Manipulation in Approval Voting
- Automatic Derivation of Results using a Computer
- 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 $f : \mathcal{P}^n \to C$, 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 several 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 the approved candidates are ranked above all the 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.

Possible Tie-Breaking Rules

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. Examples:

- The election chair may have the power to break ties.
- A designated voter may have the power to break ties.
- We may pick a winner from the set of pre-winners using a uniform probability distribution.
- We may pick a winner from the set of pre-winners using any other probability distribution.

We will try to avoid making too many assumptions as to which tie-breaking rule exactly is going to be used.

Axioms for Preferences over Sets of Candidates

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*, we assume that his preferences $\trianglelefteq$ over *sets of pre-winners* meet these axioms:

- $\trianglelefteq$ 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 $b \in B$.

Note: It isn't impossible to conceive of tie-breaking rules that don't meet these axioms (e.g. when another agent chooses the winner).

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 Optimistic Voters

We call a voter *optimistic* if his preferences over sets of pre-winners is induced only by his top candidate in each set:

$$A \trianglelefteq B \ \text{ iff } \ top(A) \preceq top(B) \quad [top(C) \in \{c^* \in C \mid \forall c \in C : c \preceq c^*\}]$$

Examples: "the election chair will break ties in my favour"; uniform tie-breaking + "extreme" utilities underlying $\preceq$.

Theorem 2 (Optimistic voters) *In approval voting, suppose that all but one voter have cast their ballot. Then, if the final voter is optimistic, he has no incentive to cast an insincere ballot.*

Proof of Theorem 2

After everyone else has voted, distinguish *pivotal*, *subpivotal* and *insignificant* candidates. The final voter has two options:

(1) Make a subset of the pivotal candidates the pre-winners.

(2) Make a subset of the subpivotal candidates together with the set of all pivotal candidates the pre-winners.

Case (1): The best option is to just approve of the most preferred pivotal candidate. Making that ballot sincere won't do any harm. ✓

Case (2a): If the most preferred pre-winner is pivotal, then our voter should have actually chosen scenario (1). ✓

Case (2b): Now suppose the top pre-winner is subpivotal, *i.e.* our voter approved of a set of subpivotal candidates. Only voting for the top subpivotal pre-winner (+ insignificant candidates above her, to make the ballot sincere) will remove some non-top pre-winners. Under optimism, this does not affect the ranking. ✓

The Case of Three Candidates

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

Theorem 3 (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 a special case of 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 3

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 4 (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`*.*

Now we would have to check a table of size 65×15:
there are $3^4 - 2^4 = 65$ situation and $2^4 - 1 = 15$ ballots.

Manual checking (though not really generation) would still be possible. But there is a better way ...

Automatic Derivation of Theorem 4

We can also automatise the checking of the dominance relations:

```
?- theorem(4).
Theorem: In approval voting with 4 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 1 conditions:
-- 4321 strictly dominates all of 431, 3.
```

For comparision, here's the output for three candidates:

```
?- theorem(3).
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 0 conditions:
```

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.
```

Implementation

The results have been generated using a program witten in Prolog. Much of it is routine, but writing the module that checks whether one set of pre-winners dominates another has been quite interesting.

This required writing a *theorem prover* for our axiom system for $\trianglelefteq$.

Here are again the axioms:

- $\trianglelefteq$ is reflexive and transitive.
- (DOM) $A \trianglelefteq B$ if $\#A = \#B$ and there exists a surjective mapping $f : A \rightarrow B$ 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 $b \in B$.

Implementing the Theorem Prover

Problem: Check whether `Bad` $\trianglelefteq$ `Good` follows from the axioms.

Idea: Treat this as a (breadth-first) *search* problem:

- `Bad` is the *initial state*; `Good` is the *goal state*.
- More precisely: any set dominated by `Good` using only (DOM) is a goal state. Note that (DOM) implies *reflexivity*.
- Axioms (ADD) and (REM) are used to *move* between states (from worse to better states).

Observe that *transitivity* is used implicitly (sequence of moves).

Note that, strictly speaking, our approach requires proof that any theorem can be derived using only a single application of (DOM) at the very end. An alternative would have been to implement (DOM) also as a move and to use just reflexivity for the goal condition (more complex).

Implementing the Theorem Prover (cont.)

Representation of states (sets of pre-winners): `[5,3,2]`, `[4]` etc.

Move from a bad to a good state (knowing the number of candidates `C`):

```
move(C, Bad, Good) :- add(C, Bad, Good) ; rem(Bad, Good).
add(C, [H|T], [X,H|T]) :- Min is H + 1, between(Min, C, X).
rem(Bad, Good) :- append(Good, [_], Bad), \+ empty(Good).
```

The goal is reached when (DOM) becomes applicable:

```
goal(CurrState, GoalState) :- dom(CurrState, GoalState).
dom([H1|Bad], [H2|Good]) :- H1 =< H2, dom(Bad, Good).
dom([], []).
```

Initiate the proof calling breadth-first search (carrying along `C`):

```
dominated(C, Bad, Good) :- solve_breadthfirst(C, Good, Bad, _).
```

The implementation of `solve_breadthfirst/4` is standard. The second argument specifies the goal; the third the initial state; the fourth would return the solution path (proof), about which we don't care here.

Implementation of Breadth-first Search

```
solve_breadthfirst(C, Goal, Node, Path) :-
  breadthfirst(C, Goal, [[Node]], RevPath),
  reverse(RevPath, Path).

breadthfirst(_, Goal, [[Node|Path]|_], [Node|Path]) :-
  goal(Node, Goal).

breadthfirst(C, Goal, [Path|Paths], SolutionPath) :-
  expand_breadthfirst(C, Path, ExpPaths),
  append(Paths, ExpPaths, NewPaths),
  breadthfirst(C, Goal, NewPaths, SolutionPath).

expand_breadthfirst(C, [Node|Path], ExpPaths) :-
  findall([NewNode,Node|Path],
          move_cyclefree(C, Path,Node,NewNode), ExpPaths).

move_cyclefree(C, Visited, Node, NextNode) :-
  move(C, Node, NextNode),
  \+ member(NextNode, Visited).
```

Examples

We can use the predicate `dominated/3` to check whether one given set of pre-winners is *definitely* dominated by another set of pre-winners according to our axioms for $\trianglelefteq$.

Here are a few examples:

```
?- dominated(5, [4,2,1], [5,3]).
Yes

?- dominated(4, [3], [4,3,2,1]).
No

?- dominated(4, [4,3,2,1], [3]).
No
```

Back to the Case of Four Candidates

Does our theorem on four candidates *matter* in practice?

The following corollary shows that it does:

Corollary 1 (Four candidates, uniform tie-breaking)
In approval voting with four candidates and 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.

Proof of Corollary 1

The reason is that the only exception to our theorem reduces to an inconsistent set of constraints under the additional assumptions of uniform tie-breaking and expected-utility maximisation.

Theorem 4 says: no incentive to manipulate *unless* our voter strictly prefers `4321` over *both* `431` and `3`.

Suppose our voter's utilities for the four candidates are u_4, u_3, u_2 and u_1, respectively, satisfying:

$$u_4 \geq u_3 \geq u_2 \geq u_1 \qquad (1)$$

So we get these constraints (comparing expected utilities):

$$\frac{1}{4} \cdot (u_4 + u_3 + u_2 + u_1) \quad > \quad \frac{1}{3} \cdot (u_4 + u_3 + u_1) \qquad (2)$$

$$\frac{1}{4} \cdot (u_4 + u_3 + u_2 + u_1) \quad > \quad u_3 \qquad (3)$$

Constraints (1)–(3) are easily seen to be inconsistent. ✓

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 btw).
  - If there are at most *three* candidates.
  - If there are at most *four* candidates, *uniform tie-breaking* is used, and all voters are *expected-utility maximisers*.
  - More?
- Next: Automatise constraint solving to see which of the exceptions for scenarios with more than four candidates matter under uniform tie-breaking with expected-utility maximisers.