Manipulation by Cloning Candidates

(with Piotr Faliszewski and Edith Elkind)

Arkadii Slinko

Department of Mathematics
The University of Auckland

COMSOC 2010 (Dusseldorf, 14 September, 2010)

Was Nader Responsible for Bush's Win?

It is widely believed that in the 2000 U.S. Green's candidate Ralph Nader have split votes away from Democratic candidate Al Gore allowing Republican candidate George W. Bush to win. The final count in Florida was:

Republican	2,912,790
Democratic	2,912,253
Green	97,488
Natural Law	2,281
Reform	17,484
Libertarian	16,415
Workers World	1,804
Constitution	1,371
Socialist	622
Socialist Workers	562
Write-in	40

Tideman's example (1987)

“ When I was 12 years old I was nominated to be treasurer of my class at school. A girl named Michelle was also nominated. I relished the prospect of being treasurer, so I made a quick calculation and nominated Michelle's best friend, Charlotte. In the ensuing election I received 13 votes, Michelle received 12, and Charlotte received 11, so I became treasurer.”

The calculation was that, being friends, Michelle and Charlotte are 'similar' and that their electorate will be split.

We would say that Tideman 'cloned' Michele.

Cloning is an important consideration for AI

Agents may need to vote on which plan to implement. Plans are tricky alternatives, so easy to clone.

Say an option

> *"Go for dinner to a particular restaurant."*

may be cloned into three:

- Go to the restaurant by a taxi;
- Go to the restaurant by a tram;
- Walk to the restaurant.

Cloning Formally

Let $R = (R_1, \ldots, R_n)$ be a profile on a set of alternatives $A = \{c, a_1, \ldots, a_m\}$. For some $k \geq 1$ we replace c with a set of alternatives $C = \{c_1, \ldots, c_k\}$ so that the new set of alternatives will be $A' = A \setminus \{c\} \cup C$.

We now extend the linear orders R_i to R'_i on A'.The new profile $R' = (R'_1, \ldots, R'_n)$ is said to be obtained by cloning from R if

$$c_s R'_i a \iff c R_i a \quad \text{for all } s \in [k] \text{ and } a \notin C,$$
$$a R'_i c_s \iff a R_i c \quad \text{for all } s \in [k] \text{ and } a \notin C.$$

Then C is said to be the set of clones of c.

Example

Here we produce three clones of a:

R_1	R_2	R_3
a	b	a
b	a	b

$\longrightarrow$

R'_1	R'_2	R'_3
a_1	b	a_2
a_2	a_2	a_3
a_3	a_1	a_1
b	a_3	b

Note that the order of clones may differ from one linear order to another.

This makes good sense since the manipulator produces clones but it is voters who determine their order.

What is a successful cloning?

We assume that voters rank clones randomly and independently so that every order on the clones is equally likely.

Definition

Given a positive real $0 < q \leq 1$, we say that the manipulation by cloning (or simply cloning) is q-successful if

(a) the manipulator's preferred candidate is not a winner of the original election, and

(b) manipulator's preferred candidate (or its clone) is a winner of the cloned election with probability at least q.

We say that cloning is 0-successful if it is q-successful for some positive (unspecified) q. This is equivalent to say that cloning would be successful if the manipulator could dictate the order of clones to each voter.

q-CLONING problem

Let $p(i, j)$ be the cost of producing jth copy of candidate c_i with $p(i, 1) = 0$. For some t we require $p(i, j) = const$ for $j \geq t$ to ensure that the price function is succinctly representable.

Definition

An instance of the q-CLONING problem for $q \in [0, 1]$ is given by the initial set of candidates $A = \{c_1, \ldots, c_m\}$, a preference profile $R = (R_1, \ldots, R_n)$ over A, a manipulator's preferred candidate $c \in A$, a parameter $t > 1$, a price function $p\colon [m] \times [t] \to \mathbb{Z}^+ \cup \{\infty\}$, a budget B, and a voting rule $\mathcal{F}$.

We ask if there exists a q-successful cloning (q-manipulation) that costs at most B.

Two special cases:

- ZERO COST (ZC): $p(i, j) = 0$ for all $i \in [m]$ and $j \in [t]$;
- UNIT COST (UC): $p(i, j) = 1$ for all i and $j \in \{2, \ldots, t\}$.

Plurality Rule

The Plurality score $Sc_P(c)$ of a candidate $c \in A$ is the number of voters that rank c first. Alternative with the largest score wins.

Theorem

For any $q < 1$, a Plurality election is q-manipulable if and only if the manipulator's preferred candidate c does not win, but is ranked first by at least one voter. Moreover, for Plurality q-CLONING can be solved in linear time. However, no election is 1-manipulable.

The idea of the proof: we clone any candidate whose Plurality score is larger than that of c:

R_1	R_2	R_3
a	c	a
c	a	c

$\longrightarrow$

R'_1	R'_2	R'_3
a_1	c	a_2
a_2	a_2	a_1
c	a_1	c

Veto (Antiplurality) Rule

The Veto score $\mathrm{Sc}_V(c)$ of a candidate $c \in A$ is the number of voters that do not rank c last.

Theorem

Any election is 1-manipulable with respect to Veto. Moreover, for Veto both 0-CLONING and 1-CLONING can be solved in linear time.

The idea of the proof: this time we clone c:

R_1	R_2	R_3
a	c	a
c	a	c

$\longrightarrow$

R'_1	R'_2	R'_3
a	c_1	a
c_3	c_2	c_1
c_2	c_3	c_3
c_1	a	c_2

Maximin (Simpson's) Rule

The Maximin score $\mathrm{Sc}_M(c)$ of a candidate $c \in A$ is the number of votes c gets in his worst pairwise contest. Winners are the alternatives with the maximal score.

Theorem

An election is 0*-manipulable by cloning with respect to Maximin if and only if the manipulator's preferred candidate c does not win, but is Pareto-optimal. Further, for Maximin* 0-CLONING *can be solved in linear time. No election is* 1*-manipulable.*

Problem

What is the supremum of such q *for which* q*-manipulable profiles with respect to Maximin exist.*

Idea of the proof

Let us consider the following profile:

$$R = \begin{array}{cccccc} 1 & 2 & 2 & 1 & 2 & 1 \\ \hline \mathrm{a} & \mathrm{a} & \mathrm{b} & \mathrm{b} & \mathrm{c} & \mathrm{c} \\ \mathrm{b} & \mathrm{c} & \mathrm{a} & \mathrm{c} & \mathrm{a} & \mathrm{b} \\ \mathrm{c} & \mathrm{b} & \mathrm{c} & \mathrm{a} & \mathrm{b} & \mathrm{a} \end{array} \longrightarrow \begin{bmatrix} 0 & 5 & 5 \\ 4 & 0 & 4 \\ 4 & 5 & 0 \end{bmatrix}.$$

a is a Condorcet winner, hence Minimax winner, and we clone it three times $a \to a_1, a_2, a_3$ arranging clones

$$R \otimes \begin{array}{ccc} 3 & 3 & 3 \\ \hline a_1 & a_2 & a_3 \\ a_2 & a_3 & a_1 \\ a_3 & a_1 & a_2 \end{array} \longrightarrow \begin{bmatrix} 0 & 6 & 3 & 5 & 5 \\ 3 & 0 & 6 & 5 & 5 \\ 3 & 6 & 0 & 5 & 5 \\ 4 & 4 & 4 & 0 & 4 \\ 4 & 4 & 4 & 5 & 0 \end{bmatrix}.$$

Now b and c are joint winners.

Borda Rule

Given an profile $R = (R_1, \ldots, R_n)$ on a set of alternatives A, the Borda score $\mathrm{Sc}_B(c)$ of a candidate $c \in A$ is given by

$$\mathrm{Sc}_B(c) = \sum_{i=1}^{n} |\{a \in A \mid cR_ia\}|.$$

Example:

3	1
a	d
c	c
b	b
d	a

$\mathrm{Sc}_B(a) = 9$
$\mathrm{Sc}_B(b) = 4$
$\mathrm{Sc}_B(c) = 8$
$\mathrm{Sc}_B(d) = 3$

Cloning $b \to b_1, b_2, b_3$ will be 1-manipulation in favour of c:

$$\mathrm{Sc}_B(a) = 15, \quad \mathrm{Sc}_B(c) = 16.$$

Problem

Characterise 1*-manipulable profiles with respect to Borda.*

Borda Rule

Theorem

An election is 0-manipulable by cloning with respect to Borda if and only if the manipulator's preferred candidate c does not win, but is Pareto-optimal. Moreover, UC 0-CLONING *for Borda can be solved in linear time.*

Idea: we simply clone c sufficiently many times.

Theorem

For Borda, q-CLONING in the general cost model is NP-hard for any $q \in [0, 1]$. Moreover, this is the case even if $p(i, j) \in \{0, 1, +\infty\}$ for all $i \in [m], j \in \mathbb{Z}^+$.

The reduction is from EXACT COVER BY 3-SETS (X3C).

Problem

What is the complexity of UC q-CLONING *for Borda for $q > 0$?*

k-Approval

For any $k \geq 1$, the k-Approval score $\mathrm{Sc}_k(c)$ of a candidate $c \in A$ is the number of voters that rank c in the top k positions.

Theorem

For any given $k \geq 2$, it is NP-hard to decide whether a given election is 0-manipulable with respect to k-Approval.

Idea: for $k = 2$ we reduce from DOMINATING SET.

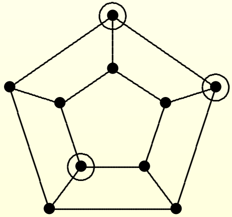

3-dominating set

The DOMINATING SET problem takes as input a graph G and a positive integer k. The question asks whether there exists a dominating set in G of size k.

k-Approval continued

Theorem

For any given $k \geq 2$, it is NP-hard to decide whether a given election is 1-manipulable with respect to k-Approval.

Problem

What is the complexity of deciding whether a given election is q-manipulable with respect to k-Approval for $q > 0$?

Copeland Rule

For simplicity we assume that the number of voters is odd.

The Copeland score $\text{Sc}_C(c)$ of a candidate $c \in A$ is the number of wins in pairwise contests against other alternatives.

Theorem

For any $q \in [0, 1]$, an election E with an odd number of voters is q-manipulable with respect to Copeland Rule if and only if the manipulator's preferred candidate c does not win, but is in the Uncovered Set of the majority relation of E. 0-CLONING can be solved in polynomial time.

Theorem

For Copeland, UC q-CLONING *is* NP*-hard for each $q \in [0, 1]$.*

We give a reduction from X3C.

Further Research

- Characterize 1-manipulable profiles for Borda.
- Can we ever manipulate Maximin elections with large probability, say $q > \frac{1}{2}$?
- What is the complexity of UC q-CLONING for Borda for $q > 0$?
- Fixed-parameter tractability analysis of NP-complete problems.
- The set of clones in any profile have hierarchical structure we call it C-structure. It makes sense to define C-structures on a finite set axiomatically An interesting question which C-structures can be realizable on profiles with single-peacked preferences.