COST Summer school on Social Choice July 2016

Lecture : Manipulation and Strategic Voting

Lecturer: Reshef Meir *Scribe: Reshef Meir*

Contents

These notes accompany a tutorial on strategic voting by Reshef Meir. This is a first draft, and may contain mistakes. Any comments regarding the content or exposition are welcome.

1 Basic Notations

We will follow the notations of [10] as much as possible.

For a finite set X, we denote by $\mathcal{R}(X)$ the set of all weak orders over X, and by $\mathcal{L}(X)$ the set of all linear (strict) orders over X. For $L \in \mathcal{L}(X)$, denote by $top(L)$ the first element of L.

A voting scenario is defined by a set of *candidates*, or *alternative*, A, a set of voters N, and a *preference profile* $\mathbf{R} = (R_1, \ldots, R_n)$, where each $R_i \in \mathcal{L}(A)$. For $a, b \in a, i \in N$, candidate a precedes b in R_i (denoted $a \succ_i b$) if voter i prefers candidate a over candidate b. Thus $top(R_i) \in A$ is i's most preferred candidate.

Definition 1.1. *A* social choice correspondence *(SCC) is a function $F : \mathcal{L}(A)^n \to 2^A \setminus \emptyset$, i.e. a function that accepts a preference profile $\mathbf{R}$ as input, and outputs a nonempty set of winning candidates.*

Definition 1.2. *An SCC F is* resolute *if $|F(\mathbf{R})| = 1$ for all $\mathbf{R}$. Resolute SCCs are also called* Social choice functions *(SCF). We typically denote SCFs by a lower case letter f.*

Definition 1.3. *A* social welfare function *(SWF) is a function $f_W : \mathcal{L}(A)^n \to \mathcal{R}(A)$, i.e. a function that accepts a preference profile $\mathbf{R}$ as input, and outputs weak order over candidates.*

We will use the term *voting rule* for a SCF unless stated otherwise.

2 Strategyproofness

Example of a manipulation Consider an election using the Borda voting rule with the following preference profile:

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

Candidate b is the winner, beating candidate a 8 points to 7. However, if voter 3 lies about his preferences and ranks candidate b last (after a, c and d), b's score goes down to 6, and a (voter 3's favorite candidate) wins! Generally, if a voter knows the voting rule being used and the preferences of the other voters (or can guess them), he may be able to bring about a more preferable result for himself by reporting a preference different from his true preference. This is called *manipulation*.

Manipulability of voting rules We would like to have a voting rule which cannot be manipulated, meaning that no voter can ever profit from lying about her preferences. But is this possible?

Definition 2.1. *A voting rule f is **strategyproof** if no (single) voter can ever benefit from lying about his preferences:*

$$\forall \mathbf{R} \in \mathcal{R}^n \ \ \forall i \in N \ \ \forall R'_i \in \mathcal{R} : f(R'_i, \mathbf{R}_{-i}) \preceq_i f(\mathbf{R})$$

For facility location and classification the definition is analogous, except that preferences are locations or binary label vectors rather than orders.

Definition 2.2. *A voting rule is **manipulable** if it is not strategy-proof.*

Claim 2.3. *If there are exactly two candidates then the Majority voting rule is strategy-proof.*

Proof. Denote $A = \{a, b\}$ and assume WLOG that candidate a was selected by the Majority voting rule (for some given preference profile $\mathbf{R}$). Let $i \in N$ be one of the voters. If $b \prec_i a$, then voter i's favorite candidate has already won, and she certainly has nothing to gain by lying about her preferences. On the other hand, if $a \prec_i b$, then changing voter i's preference so that b is ranked below a only lowers b's score and cannot possibly cause b to win. In either case, voter i cannot benefit from lying about her preferences. □

2.1 The Gibbard-Satterthwaite theorem

Recall the definitions of dictatorial rules and duples:

A voting rule f is **dictatorial** if there is an individual (the dictator) whose most preferred candidate is always chosen by f:

$$\exists i \in N \quad \forall \mathbf{R} \in \mathcal{L}(A)^n : f(\mathbf{R}) = top(R_i).$$

A voting rule f is a **duple** if there are only two possible winners: $|\{f(\mathbf{R}) : \mathbf{R} \in \mathcal{L}(A)^n\}| \leq 2$.

Theorem 2.4. *(The Gibbard-Satterthwaite theorem [26, 71]) If there are at least three candidates, any deterministic voting rule is strategy-proof if and only if it is either dictatorial or a duple.*

It is easy to see that a dictatorial rule is SP, since the dictator is always best off reporting the truth and all other voters have no power; if we allow duples, we can arbitrarily select two candidates a, b and hold a majority vote between them.

There are many different proofs of the Gibbard-Satterthwaite theorem. Several simple proofs can be found in [73]. In the appendix bring one of these proofs, under the simplifying assumption that $n = 2$.

2.2 Manipulations in other domains

Facility location Facility location can be thought of as a special case of voting, where the alternatives A are possible locations for a facility in some metric space. Each agent is assumed to prefer the facility to be as close as possible to her location, thus instead of reporting her full preference order R_i, she only needs to report her location (say some point $x_i \in \mathbb{R}^d$ or on a graph G).

While there are many possible aggregation rules, the most natural one would simply take all of the locations $x_1, \ldots, x_n$ and return the point $a \in A$ that minimizes the average distance to all input points. It is not hard to see that an agent i can manipulate by reporting a different location x'_i (typically more extreme), thereby moving the facility closer to her real location.

Binary classification A classification setting consists of an input space and a class of classifiers (known as the *concept class*), that is, functions from the input space to the set of *labels* $\{+, -\}$. A classification mechanism receives as input a *dataset*—a set of input points and their labels, where each such labeled input point is known as an *example*—and must return a classifier from the given class that classifies the given examples as well as possible. For instance, the input space might be images (represented as matrices of pixels), and the dataset might label different images as showing a human face (positive label) or not showing a face (negative label).

In our setting the labels are reported by strategic agents. Each agent controls a subset of the dataset, i.e., its own subset of (labeled) examples. The input points controlled by each agent are known, but their labels are private information. Given the reported labels, the classification mechanism selects a classifier. However, by lying the agents may achieve a classifier that better reflects their own labels, at the expense of overall accuracy.

Judgement aggregation In judgement aggregation problems there is a finite agenda of k binary issues, with a logical constraint on the valid combinations $V \subseteq \{0,1\}^k$ (e.g. "$C = (A \;\; or \;\; B)$" means that the assignment $A = 0, B = 0, C = 1$ is not valid). An aggregation function collects the opinions of all judges (i.e. valid binary vectors $x_i \in V$), and outputs a valid outcome. In contrast to facility location, there is typically no optimization goal, but rather we want aggregators that have certain axiomatic properties like IIA ($f(x_1, \ldots, x_n)_j = f_j(x_{1j}, \ldots, x_{nj})$ for all $j \leq k$), unanimity ($f(x, \ldots, x) = x$) and so on.

Since there is no one clear notion of utility, manipulation is also not uniquely defined [18]. For example we can think of "optimistic" manipulation, where the agent succeeds in changing the outcome on one particular issue y_j to x_{ij}, regardless of how other issues are affected; "pessimistic," where the outcome must be at least as good for i on every issue j; and "utilitarian," where the agent cares about the total number of issues that match her opinion. Under the latter notion of agents' utility, JA is a special case of facility location, where every agenda corresponds to a subset of $\{0,1\}^k$. For example, for the agenda in our example above, every opinion x_i is a point in $A = \{(000), (010), (100), (111)\}$.

The exact formulation of strategyproofness may differ according to the domain. See definitions and discussion in the mentioned papers, and in references therein: Facility location [63, 19]; Binary Classification [44]; Judgement Aggregation [35, 18].

3 Robustness of the G-S Theorem

There have been many attempts to test whether the strong negative result of G-S can be mitigated by slightly relaxing our requirements or assumptions. In this section we consider two approaches that actually show the robustness of the negative result:

1. A lower bound for the frequency of manipulation.

2. An impossibility theorem for randomized rules.

3.1 Frequency of manipulation

While the G-S theorem shows that for every rule there is *some* manipulation, perhaps there are rules for which those manipulations are too rare for us to care about. If, for example, there is only one or two "dangerous" profiles out of $(m!)^n$ then maybe we can live with such "almost-SP" rules.

Unfortunately, it can be shown that in every reasonable voting rule there will be much more manipulations than that. More formally, we define the *manipulation power* of voter i in rule f as

$$M_i(f) = \frac{1}{(m!)^n} \sum_{\mathbf{R} \in \mathcal{L}(A)^n} [\![\exists R'_i \text{ s.t. } f(\mathbf{R}_{-i}, R'_i) \succ_i f(\mathbf{R})]\!] ,$$

i.e. as the probability that in a random profile (sampled uniformly according to the impartial culture) voter i will have some beneficial manipulation.

Note that we can rephrase the G-S theorem as follows:

Theorem 3.1. *For deterministic voting rule f, either f is dictatorial, or a duple, or $\sum_{i \in N} M_i(f) > 0$.*

Now that we have a quantitative measure for manipulation, we would like to similarly relax the notions of dictatorships and duples.

Definition 3.2. *We say that a rule f is ε-bad if there is some other voting rule g s.t.*

1. *g is either a duple or dictatorial.*

2. *$Pr_{\mathbf{R} \sim U(\mathcal{L}(A)^n)}(f(\mathbf{R}) = g(\mathbf{R})) \geq 1 - \varepsilon$, i.e., f and g agree on all input profiles except ε.*

In particular, dictatorial function and duples are 0-bad. The question of whether we can lower bound the manipulation power of ε-bad rules for $\varepsilon > 0$ is a relatively recent one.

The first studies of similar questions [15, 62] were more algorithmic in nature, and provided algorithms that w.h.p can tell whether a manipulation exists or not. Friedgut et al. [25] was the first to prove such a lower bound on the manipulation power (under a slightly different definition), albeit only for Neutral rules with $m = 3$.

The state-of-the-art result now shows that:

Theorem 3.3 (Mossel and Rácz [49])**.** *For any $\varepsilon \geq 0$, and any deterministic voting rule f, either f is ε-bad, or $\sum M_i(f) > poly(\frac{1}{n}, \frac{1}{m}, \varepsilon)$.*

That is, to keep a low probability of manipulation, we must select a rule that is very close to being a duple or dictatorial. Note e.g. that all anonymous and unanimous rules (including all scoring rules, Copeland, Kemeny, etc.) are very far from being "bad" and thus have at least a polynomial probability of manipulation.

Also note that we get the Gibbard-Satterthwaite theorem (Theorem 3.1) as a special case for $\varepsilon = 0$.

Group manipulations One may claim that the result above does not really convince that manipulations are common. Indeed, a polynomial fraction may still be quite small. Even in the Plurality rule, which is the most manipulable, the probability of manipulation is about $\frac{1}{\sqrt{n}}$. When there are millions of voters, this becomes negligible.

However, things look less rosy when we consider groups of voters that can manipulate together. If we allow sufficiently large groups, then we will almost always have some groups (or even many groups) who can get a better outcome for all voters in the group by manipulating together.

Formally, f has a *group manipulation* under profile $\mathbf{R}$ if there is a set of voters $\mathbf{S} \subseteq N$, and $\mathbf{R}'_S$ s.t. for all $i \in S, f(\mathbf{R}_{-S}, \mathbf{R}_S) \succ_i f(\mathbf{R})$. f is *group strategyproof* if there is no group manipulation under any profile.

In fact, there is a tight connection between the average-case manipulability, and the minimal size of the coalition that is required to change the outcome. For most common voting rules, coalitions of size at least $\sim \sqrt{n}$ can decide the identity of the winner in almost every profile [77]; for smaller coalitions the probability they have any effect on the outcome is small.[1]

3.2 Randomized voting rules

Suppose we allow our voting rule to flip coins, i.e. return different outcomes with certain probabilities. More formally, a randomized voting rule is is a function that maps any profile $\mathbf{R}$ to a lottery (probability distribution) $\sigma \in \Delta(A)$. We denote by $f(\mathbf{R}, c) = \sigma(c) \in [0, 1]$ the probability that f returns $c \in A$ on profile $\mathbf{R}$. Note that we can also think about f as a lottery over deterministic voting rules.

It is easy to see that we can find such rules that violate the Gibbard-Satterthwaite conditions. For example, we can think of a rule that return any candidate with equal probability, regardless of the profile.

We can also consider somewhat more sophisticated rules, for example, selecting one voter at random, and then use this voter as a dictator. Or, alternatively, select two candidates at random, and then hold a Majority vote among them.

All of these rules are *strategy-proof ex-post*, that is, no voter has a manipulation even after we reveal the outcome. We can also think of more refine notion of manipulation, where a voter gains *in expectation* by deviating. To define expected utilities, we need to extend the voter preferences to *cardinal utilities*, i.e. assume that each voter i assigns to each candidate c some numerical utility $u_i(c) \in \mathbb{R}$.[2]

Strategyproofness means that in any cardinal profile $\mathbf{u} = (u_i(c))_{i \in N, c \in A}$, any voter $j \in N$, and any alternative utility scale u'_j,

$$E_{c \sim f(\mathbf{u})}[u_j(c)] \geq E_{c \sim f(\mathbf{u}_{-j}, u'_j)}[u_j(c)].$$

It may seem that we might find rules that are SP in expectation, but not ex-post. In turns out however, that the options we suggested above are roughly the *only* randomized rules that are strategyproof, even in expectation.

[1]By small here we mean “goes to 0 as n grows. However in the part above we considered such probabilities (polynomial rather than exponential) as large. All is relative in life.

[2]We typically assume that $u_i(c) \neq u_i(c')$ for all $c, c' \in a$.

Definition 3.4. *A voting rule is* unilateral *if there is some $j \in N$ s.t. for all $u_j, \mathbf{u}_{-j}, \mathbf{u}'_{-j}$, $f(\mathbf{u}_{-j}, u_j) = f(\mathbf{u}'_{-j}, u_j)$.*

That is, the outcome is set exclusively based on the input from j.

It is not hard to see that among the unilateral rules, some are strategyproof (e.g. dictatorial rules), while others are not (e.g. selecting j's top choice w.p. 1/3, and j's last choice w.p. 2/3).

Theorem 3.5. *[27] A (randomized) voting rule f is strategyproof in expectation, if and only if it is a lottery over (deterministic) duples and strategyproof unilateral rules.*

In other words, allowing randomization extends the set of SP voting rules, but still doesn't allow for "reasonable" rules. In the next section, we will discuss several other approaches to guarantee truthfulness.

4 Achieving truthfulness in voting

We will focus on four approaches, each of which attains truthfulness under certain assumptions:

Subsection	Approach	Assumptions
1	Domain restriction	The preferences $\mathbf{R}$ have a certain structure
2	Complexity barriers	Voters have bounded computational resources
3	Approximation	-
4	Differential privacy	Voters will only manipulate for a substantial gain
5	Payments	We can charge arbitrary amounts of money from voters, Voters have quasi-linear utility for money

The two latter approaches also assume voters have cardinal utilities for candidates.

4.1 Domain restriction

Suppose voters are voting on where to place a public library along a street. Naturally, each voter prefers the library to be located as close as possible to her house (whose location is private information). If we number all the addresses in the street from (say) east to west, then it is not possible for example that a voter prefers the east-most location, then the west-most location, and then some location in the middle. Thus not every preference profile is possible.

More formally, given an order L on candidates A, a preference order R_i is *single peaked* w.r.t. L if there is some "peak candidate" a^*_i s.t. for all $x, y \in A$, if $L(x)$ is between $L(a^*_i)$ (included) and $L(y)$ (excluded), then $a^*_i \succeq_i x \succ_i y$. In particular, a^*_i is i's most preferred candidate.[3]

Example: suppose $A = \{a_j\}_{j=1}^5$ are ordered from a_1 to a_5, i.e. $L(a_j) = j$. Then the following orders are SP w.r.t. L:

- $a_1 \succ a_2 \succ a_3 \succ a_4 \succ a_5$ (a_1 is the peak).
- $a_3 \succ a_2 \succ a_4 \succ a_5 \succ a_1$ (a_3 is the peak).
- $a_4 \succ a_3 \succ a_2 \succ a_1 \succ a_5$ a(a_1 is the peak).

On the other hand, $a_2 \succ a_1 \succ a_3 \succ a_5 \succ a_4$ is not SP w.r.t. L, since for $x = a_4, y = a_5$ we get a violation of the definition.

[3]Note that this model is highly related to the facility location problem we mentioned earlier. However here we do not assume any metric or cardinal utilities.

The Median Voter rule Given a linear order L over alternatives A, and a preference profile $\mathbf{R}$ that is single-peaked on L:

1. Ask each voter $i \in N$ to report her peak $a_i = top(R_i)$.
2. Set $l_i = L(a_i)$, i.e. the order of a_i in L.
3. Find $j^* \in N$ s.t. $k_{j^*} \in \text{median}\{l_1, \ldots, l_n\}$ (if there are more than one, select arbitrarily).
4. Return $f(\mathbf{R}) = a_{j^*}$.

Theorem 4.1 (Moulin [50]). *The Median Voter rule is group-strategyproof (in particular strategyproof).*

Proof. Assume towards a contradiction that there is a subset $S \subseteq N$ that can manipulate by reporting $\mathbf{R}'_S$, and denote $c^* = f(\mathbf{R}), c' = f(\mathbf{R}_{-S}, \mathbf{R}'_S)$. Clearly S does not contain any voter i s.t. $a_i = f(\mathbf{R})$ as such a voter strictly loses from any manipulation. Denote by $S_R, S_L \subseteq S$ the sets of agents whose peaks are to the right and to the left of $f(\mathbf{R})$, respectively. W.l.o.g. $L(c') < L(c^*)$, i.e. c' is to the left of c^*. It cannot be that S_R is empty, since no report by S_L can move the median left. However, consider some $i \in S_R$, then $L(c') < L(c^*) < L(a_i)$ thus i strictly loses from this manipulation. A contradiction. □

It also holds that the Median Rule is IIA, and that under SP preferences the Condorcet winner always exists, and is always the median.

4.2 Complexity barriers

[4] Another approach is to look for voting rules where even if manipulations exist, it might be computationally hard for the voter (the potential manipulator) to actually find what to report in order to promote her preferred candidate. Note that since the number of possible reports is $m!$, a brute-force search is typically infeasible.

MANIPULATION$_\mathbf{f}$: given a set of candidates A, a group of voters N, a voting rule f, a manipulator $i \in N$, a preference profile of all voters except i $\mathbf{R}^N_{-i} = (R_1, \ldots, R_{i-1}, R_{i+1}, \ldots R_n)$, and a specific candidate $p \in A$: Answer whether the manipulator can provide a preference R^*_i such that $f(\mathbf{R}_{-i}, R^*_i) = p$.

Notice the following observations:

- The problem is a decision problem. We would like to know if there exists such a preference.
- The problem is defined more strongly than the simple definition of manipulation. The question is not whether the manipulator can guarantee a victory of a more preferred candidate, but whether he can guarantee the victory of a specific candidate.
- The Gibbard-Satterthwaite theorem specifies that one of the voters can manipulate the results. We don't know if this is true under the given circumstances, where i and p were given.

[4] This subsection is based on notes by Ariel Procaccia, Ezra Resnick & Ariel Imber.

A greedy algorithm This algorithm generates the preference of the voter i such that the candidate p will be chosen, or returns that this is not possible. The algorithm does the following:

1. Rank p in 1^{st} place.
2. While there are candidates that were not ranked yet:
   - If there exists a candidate that can be ranked in the next spot without preventing p from winning, rank that candidate in the next spot.
   - otherwise - declare that the desired preference does not exist.

Example

In this example m=4, n=3, and the **Borda** voting rule (where every voter gives m-1 points to his top candidate, m-2 to the next, etc.) will be used. The preference of voters 1 and 2 is a, p, b, c and the result of the algorithm for the 3^{rd} voter is:

1. Rank p in 1^{st} place, and find that with the two other voters p has 7 points.
2. Notice that a can't be ranked next, since he'll end with 8 points. Rank b in 2^{nd} place, leaving him with 4 points.
3. Notice that a can't be ranked next, since he'll end with 7 points. Rank c in 3rd place, leaving him with 1 point.
4. Rank a in 4^{th} place, leaving him with 6 points.

According to f, p wins with 7 points (while a, b, and c have 6, 4, and 1 points respectively).

Proposition 4.2. *Let f be any positional scoring rule, then the greedy algorithm always decides MANIPULATION$_f$ correctly.*

In fact, the algorithm works for a much broader class of voting rules that also includes Copeland and Maximin [6], but we only prove here a restricted version.

Proof. It's obvious that if the algorithm found a preference resulting in p winning, the decision that the manipulator can provide a preference such that p wins is correct. Let's examine the case where the algorithm could not find such a preference.

Suppose that the algorithm stopped ranking when the subset $U \subseteq A$ were not yet ranked, and assume that there exists a preference $R^{'}$ for the manipulator that causes p to win. Denote $u \in U$ the top candidate of the subset U according to $R^{'}$. Complete the partial preference created by the algorithm to a preference R^* that u is ranked in the 1^{st} place that was not already ranked by the algorithm.

(Example: Let $A = \{a, b, c, d, p\}$, suppose that the algorithm generated the partial preference (p, c, d) and assume that $R^{'} = (b, c, d, p, a)$. Then $U = \{a, b\}$ and $u = b$.)

Let $\mathbf{R}' = (\mathbf{R}_{-i}, R'), \mathbf{R}^* = (\mathbf{R}_{-i}, R^*)$. Since R' is a successful manipulation, $s(\mathbf{R}', p) > s(\mathbf{R}', x)$ for all $x \neq p$ (taking tie-breaking into account).

For all $x \in A \setminus (U \cup \{p\})$, we know that $s(\mathbf{R}^*, p) > s(\mathbf{R}^*, x)$ since otherwise the algorithm would have stopped earlier. For all $x \in U$, x is ranked (weakly) lower in R^* than in R', whereas p is pushed up. Thus

$$s(\mathbf{R}^*, p) \geq s(\mathbf{R}', p) > s(\mathbf{R}', x) \geq s(\mathbf{R}^*, x).$$

Thus p must win in $\mathbf{R}^*$ as well, as required. □

Voting rules that are NP-hard to manipulate Note first that in some rules, such as Kemeny, it is already NP-hard to compute the outcome of a given profile. Such rules are clearly hard to manipulate but this is not an interesting results.

Theorem 4.3 ([6]). *There is a voting rule f such that: I) $f(\mathbf{R})$ can be computed in polynomial time; II)* $\mathbf{MANIPULATION_f}$ *is an NP-Complete problem.*

The original proof in [6] used a variation of Copeland, however similar hardness results hold for common voting rules such as STV [5]. We provide a hardness-of-manipulation proof for another rule called Plurality with Preround in the appendix.

Few candidates and Coalitional manipulations Note that for a fixed number of candidates m, there is a trivial polynomial-time algorithm for computing a manipulation: simply try all $m!$ possible ballots, which is also a fixed number.

The *coalitional manipulation* problem is similar to $\mathbf{MANIPULATION_f}$, except it gets as input a set of manipulators I rather than a single manipulator. The number of possible combined ballots that I can submit is $(m!)^n$, so even for a fixed m this is a nontrivial problem. Indeed, for many voting rules the coalitional manipulation problem turns out to be NP-hard even for small values of m [14].

A particular problem of interest that has been open for several years was the complexity of Borda manipulation of $n = 2$ voters (and any number of candidates). This was shown in be an NP-hard problem in [17, 8]. For more details on complexity barriers see [10], Section 6.4.

4.3 Approximate scoring rules

Consider a rule with a natural notion of candidate score $s(a)$, such as Plurality, Borda, or Maximin. For any outcome alternative $a \in A$ we can compare $s(a)$ to the score of the "true" winner of the voting rule $s(f(\mathbf{R}))$. Thus a (randomized) voting rule f is a γ output-approximation of a scoring rule g if for any profile $\mathbf{R}$, $E[s_g(f(\mathbf{R}))] \geq \gamma \cdot s_g(g(\mathbf{R}))$, where s_g is the scoring used by rule g.

Procaccia [61] suggested to limit ourselves to voting rules that are strategyproof (i.e., fall under Gibbard [27] characterization), and within this set select a rule that approximates our target rule g as well as possible.

For PSRs, this is done by selecting one voter uniformly at random, and define a particular lottery over unilateral rules, that depends on the scoring function α (e.g. for Plurality, this always selects the most preferred candidate of this voter). He shows the following:

Theorem 4.4 (Procaccia [61]). *Let g be any PSR, then there is a randomized strategyproof voting rule f such that f is a $\Omega(\frac{1}{\sqrt{m}})$ output-approximation of g. This bound is tight for the Plurality rule.*

That is, for a fixed number of candidates m, we get a fixed approximation ratio regardless of the number of voters.

Notably, this approach requires the weakest assumptions on voters' behavior - we only assume they follow their weakly dominant strategy to be truthful. On the other hand, we get an outcome that is only an approximation. Further, this is approximation in expectation, and in a particular realization the outcome may be much worse than g.

Approximation with almost-SP rules Núnez and Pivato [54] suggested to use strategy rules as Procaccia did, but instead using the SP rule (say, random dictator) all the time, use a random dictatorial rule with some low probability ε, and run the target rule g with probability $1-\varepsilon$. They show that the combined rule guarantees a good approximation of the target rule g. In addition, while the combined rule is not strategyproof, under mild assumptions on the preferences of the population they show that as the population grows the probability that a voter will want to manipulate goes to 0.

4.4 Differential privacy

Consider a randomized algorithm f that operates on some data $\mathbf{R}$ and returns some outcome $a \in A$. We can think of $\mathbf{R}$ as a table with one row per person, and where columns specify various information about this person.

An algorithm f is *differentially private* if for all , that differ by 1 person, and all $S \subseteq A$ we have $Pr[f(\mathbf{R}) \in] \cong Pr[f() \in]$. That is, we cannot tell by the outcome (almost) anything about a particular person in the data.

There is abundant literature on constructing differentially private algorithms, or modifying existing algorithms so as to make them differentially private. This is typically by adding noise to the dataset in a particular way, see e.g. [21].

Differentially private voting rules Intuitively, a (randomized) voting rule f is differentially private if for all profiles $\mathbf{R}, \mathbf{R}'$ that differ by a single voter, and for any subset of alternatives $S \subseteq A$, $Pr[f(\mathbf{R}) \in S] \cong Pr[f(\mathbf{R}) \in S]$. That is, a single voter has very small influence on the outcome. While it seems that this is the case in most voting rules, note that in standard deterministic rules (like Plurality), there are very few profiles in which a voter is pivotal, but in those profiles the voter has large influence (changes the outcome with certainty). In contrast, in differentially private rules, in *every* profile a single voter has negligible influence.

How does this relate to strategyproofness? Since the influence of every voter is small, the (expected) gain from manipulation is also small. Thus the main idea is to take some deterministic voting rule f, and modify it by adding some noise to the votes before computing the outcome. If done properly, this will have very little effect on the outcome in most profiles, while reducing the incentive of voters to manipulate.

Approximately SP voting

Definition 4.5. *Fix a pseudo-metric $d : \mathcal{L}(A) \times \mathcal{L}(A) \to \mathbb{R}_+$. A randomized voting rule f is a δ input-approximation of another voting rule g, if for any profile $\mathbf{R}_g$ and every coin toss, there is another profile $\mathbf{R}_f$ such that:*

- $d(\mathbf{R}_f, \mathbf{R}_g) \leq \delta$,
- $f(\mathbf{R}_f) = g(\mathbf{R}_g)$.

E.g., if d counts the number of voters who vote differently, then for every $\mathbf{R}$ every realized outcome of f we can change the vote of at most δ voters so as to get the "correct" outcome $g(\mathbf{R})$. Note that this is a very different notion of approximation than the one used in Section 4.3. As a more concrete example, consider the Plurality rule g, and define f by first sampling k random voters and replace their vote with an arbitrary candidate, and then compute the Plurality winner. Then f is a k-approximation of Plurality.

Theorem 4.6 (Birrel and Pass [9])**.** *For any deterministic voting rule g over m candidates, and any $\varepsilon > 0$, there is a randomized voting rule f that is:*

- *ε-strategyproof, i.e. no voter has a probability greater that ε to improve her utility by manipulation;*
- *δ input-approximation of g, for $\delta = O(\frac{m^2}{\varepsilon})$.*

In other words, we can construct a voting rule f where the incentive to manipulate is as low as we want. This voting rule will coincide with g except in "borderline" profiles where a small number of voters could change the outcome. The size of this margin is inversely proportional to the incentive to manipulate.

While this result (and other, more nuanced results in the paper) is mathematically sophisticated, it is not clear to what extent they are required in practice. In reality, voters have uncertainty over the outcome even without adding noise, and thus they are already "differentially private" to some extent from the perspective of the voter. Yet, voters do manipulate even though their expected gain is negligible, which raises the question if ε-strategyproofness is a good enough guarantee.

4.5 Payments

Adding payments allows us to transfer utility between agents with much flexibility, thereby aligning their incentives. The main tool we apply is a powerful mechanism with many application in mechanism design, called VCG.

Voting with money To be able to compare voting outcomes and money, we assume voters have cardinal utilities over candidates as in Section 3.2. That is, $u_j(c) \in \mathbb{R}$ is the utility that voter j assigns to candidate c. We now reformulate the voting process as an auction, where each voter announces her value for each candidate. We can think of these values as "bids", ad a voter may bid higher or lower than her true value. Each agent will be charged some price p_j, and her utility will be $u_j(c) - p_j$ where c is the elected candidate (quasi-linear utilities).

The logical thing to due given the reported values, is to select the candidate that maximizes the social welfare, i.e. $c^* = \text{argmax}_{c \in A} \sum_{i \in N} u_i(c)$. We then need to decide the amount we charge each agent.

As a first attempt, we can charge each voter j the amount of her bid, i.e., $p_j = u_j(c^*)$. However, it is not hard to see that this is not strategyproof. Indeed, suppose all voter bid 1 on c^* and 0 on all other candidates. Then j can manipulate by reporting $u'_j(c^*) = 0$: c^* will still be selected, the payment of j will drop, and her utility will strictly increase.

VCG The Vickrey-Clarke-Groves mechanism collect all information from the agents. Then it computes the optimal outcome, and charges each agent i the "damage" that this agent inflicts upon the other agents. That is, the difference between the maximal social welfare in a world where i does not exist, and the maximal social welfare (of all except i) in the current world. For a thorough exposition of VCG see [53].

While the VCG mechanism guarantees truthfulness in a wide variety of applications, it turns out that as a voting rule it has a particularly attractive form:

1. Collect $u_i(c)$ for all $i \in N, c \in A$.
2. Let $c^* = \text{argmax}_{c \in A} \sum_{i \in N} u_i(c)$.
3. For each $j \in N$:
   (a) Compute $SW_{N \setminus \{j\}} = \max_{c \in A} \sum_{i \in N \setminus \{j\}} u_i(c)$ (the social welfare when j does not exist).
   (b) Compute $SW_N^{-j} = \sum_{i \in N \setminus \{j\}} u_i(c^*)$ (the current social welfare without j).
   (c) Set $p_j = SW_{N \setminus \{j\}} - SW_N^{-j}$.

In other words, agent j only pays *if she is pivotal*, and in that case only pays the minimal bid that would still make c^* win. By definition, VCG always selects the candidate maximizing the social welfare.

Theorem 4.7 (Clarke [13])**.** *The VCG voting rule is strategyproof. That is, for any* $\mathbf{u} \in \mathbb{R}^{n \times m}$

For a full proof see [53]. However the intuition is simple: a voter j can try to manipulate in one of two ways. She can try to reduce her bid on c^*, or increase her bid on another candidate c' to make it win. In the former case, p_j (and thus the utility of j) will not change unless the winner changes. But if the winner changes to j', then $u_j(c^*) \geq u_j(c') + p_j$, which means j does not gain. Likewise, in the latter case, in order to change the winner to c', j must in crease her bid enough so that $u_j(c') + p'_j \geq u_j(c^*)$.

Thus VCG supposedly solves the incentives problem completely, but under fairly strong assumptions: voters have cardinal preferences over candidates, we can charge voters arbitrary amounts of money, and voters attribute quasi-linear utility to money. A recent paper shows that the last assumption cannot be relaxed: under mild requirements, if voters have non-quasi linear utilities then there are no strategyproof rules even with money []

5 Voting Equilibrium

One of the earliest formal analyses of strategic voting is due to Farquharson [23], which models it as an extensive-form game . Notably, this brilliantly text precedes even the G-S theorem and is highly recommended for reading. However we will take a different approach that follows the more common framework of normal-form games.

g_1	a	b	c
a	a	a	a
b	b	b	b
c	c	c	c

g_2	a	b	c
a	a	a	a
b	a	b	b
c	a	b	c

g_3	x	y
a	a	b
b	b	c
c	c	a

g_4	x	y	z	w
a	ax	ay	az	aw
b	bx	by	bz	bw
c	cx	cy	cz	cw

Figure 1: Four examples of game forms with two agents. $g_1 = f_1$ is a dictatorial game form with 3 candidates (the row agent is the dictator). $g_2 = f_2$ is the Plurality voting rule with 3 candidates and lexicographic tie-breaking. f_3 and f_4 are non-standard game forms. In g_3, $A_1 = C = \{a,b,c\}, A_2 = \{x,y\}$. Note that g_4 is completely general (there are 3×4 possible outcomes in C, one for each voting profile) and can represent any 3-by-4 game.

$(f_2, \mathbf{R}^1)$	a	b	c
a	$3,1*$	$3,1*$	$3,1*$
b	$3,1$	$2,3$	$2,3$
c	$3,1$	$2,3$	$1,2$

$(f_2, \mathbf{R}^2)$	a	b	c
a	$1,2$	$1,2$	$1,2$
b	$1,2$	$2,3$	$2,3$
c	$1,2$	$2,3$	$3,2$

Figure 2: Consider first $\mathbf{R}^1 = (a \succ_1 b \succ_1 c, b \succ_1 c \succ_1 a)$. The ordinal utilities of both voters (higher is better) are on the left figure, and Nash equilibria are marked with *. We show the same in the right figure for the same game form with preferences $\mathbf{R}^2 = (c \succ_1 b \succ_1 a, b \succ_2 c \succ_1 a)$.

Voting rules, game forms, and games A game form g allows each agent $i \in N$ to select an action a_i from a set of messages A_i. Thus the input to g is a vector $\mathbf{a} = (a_1, \ldots, a_n)$ called an *action profile*. Then, g chooses a winning alternative—i.e., it is a function $g : \mathcal{A} \to C$, where $\mathcal{A} = \times_{i \in N} A_i$, and C is some set. See Fig. 1 for examples.

A *game* is attained by adding either cardinal or ordinal utility to a game form. The linear order relation $R_i \in \mathcal{L}(C)$ reflects the preferences of agent i. That is, i prefers c over c' (denoted $c \succ_i c'$) if $(c, c') \in R_i$.

A game form g is *standard* if $A_i = \overline{A}$ for all i, and $\overline{A}$ is either $\mathcal{L}(A)$ (the set of permutations over A) or a coarsening of $\mathcal{L}(A)$.

Thus, a standard game form coincides with our previous definition of a voting rule, or SCF. Each particular voting instance that includes voters' real preferences is a game. Indeed, most common voting rules including SCRs, Maximin, Kemeny etc. are standard game forms once we fix a deterministic tie-breaking rule.

Approval is an example of a non-standard game form (note that it does not match our definition of SCF or SCC). Another example is the Direct Kingmaker rule, where voter n specifies a number $i \in \{1, \ldots, n-1\}$, and the winner is *top*(R_i).

Nash equilibrium Once we have a game, the game theoretic approach is to analyze its equilibria. Figure 2 shows two different games derived from the Plurality voting rule, and their pure Nash equilibria.

Definition 5.1. *Given a game* $(f, \mathbf{R})$*, an action profile* $\mathbf{a}$ *is a* pure Nash equilibrium *if for all* $i \in N$ *and all* a'_i *it holds that* $f(\mathbf{a}) \succeq_i f(\mathbf{a}_{-i}, a'_i)$.

A *mechanism* can be any game form g whose range is A.

Denote by $NE_g(\mathbf{R}) \subseteq A$ all candidates that win in some pure Nash equilibrium of the game $(g, \mathbf{R})$.

Consider the Plurality voting with $n \geq 3$ voters. It is easy to see that any profile in which all voters vote for the same candidate is a Nash equilibrium, i.e. $NE_{\text{Plurality}}(\mathbf{R}) = A$ for all $\mathbf{R}$. This is true even if this candidate is ranked last by all voters in $\mathbf{R}$, since no single voter can change the outcome. This already indicates that using Nash equilibrium as way to predict or recommend voters' actions is problematic. Indeed, in all common voting rules, almost any voting profile is a Nash equilibrium regardless of preferences. In the rest of this section we will consider refinements and restrictions that aim to "narrow" the set of Nash equilibria and make it more useful as a solution concept.

5.1 Implementation

Consider any (non-standard) SCC F, i.e. a function that maps preference profiles to subsets of outcomes.

Implementation of F by a mechanism g means that given any profile $\mathbf{R}$, all "reasonable" outcomes of g when voters have preferences $\mathbf{R}$ are $F(\mathbf{R})$. Reasonable in that sense typically means a common solution concept such as Nash or strong equilibrium.

A most natural question is which voting rules implement themselves, if any. This question can be extended by allowing arbitrary mechanisms that are not necessarily voting rules, and ask if a voting rule can be implemented by such a mechanism.

NE implementation Formally, a mechanism g implements a SCC F in NE, if $NE_g(\mathbf{R}) = F(\mathbf{R})$ for all $\mathbf{R}$.

Suppose g is a dictatorship. That is, for some $i \in N$, $g(\mathbf{R}) = top(R_i)$ for all $\mathbf{R}$. In every Nash equilibrium of g, i ranks $top(R_i)$ first, and thus $g(\mathbf{R}) = top(R_i)$. Therefore, A dictatorship g (trivially) implements itself in NE.

But are there less trivial examples of such voting rules? The answer is quite negative, even if we allow the designer to use arbitrary mechanisms.

Theorem 5.2 (Maskin [36])**.** *No voting rule except dictatorships and duples can be implemented in NE by any mechanism. In particular no manipulable voting rule implements itself in NE.*

Maskin further showed that if we want to implement SCCs rather than SCFs (i.e. rules that allow for more than one winner) then results are more positive. A trivial example is the SCC $F(\mathbf{R}) = A$, which can clearly be implemented (e.g. by Plurality with $n \geq 3$ voters).

A less trivial example is $F(\mathbf{R})$ which returns all Pareto-optimal outcomes of $\mathbf{R}$. Maskin defines a certain property of SCCs that he calls *monotonicity*.[5] Formally, that if $a \in F(\mathbf{R})$ and in $\mathbf{R}'$ all voters only move a up their ranking (possibly changing order between other candidates), then $a \in F(\mathbf{R}')$. *No veto power* is a slight relaxation of unanimity, stating that if $n-1$ rank $a \in A$ at the top then $a \in F(\mathbf{R})$.

Theorem 5.3 (Maskin [36])**.** *Consider an SCC F with $n \geq 3$ voters. F can be implemented in NE if it is monotone and has no veto power.*

[5]This property is widely known as *Maskin monotonocity*, to distinct, and it is a close reminiscent of IIA. Note that there is also a weaker notion of motonicity in the literature, where voters move up a without changing the order among other candidates. Most common voting rules are monotone under the second definition.

On the other hand, a slight relaxation of these conditions results in non-implementable SCCs. The main idea of Maskin's proof is to construct a mechanism where each voter reports all of $\mathbf{R}$, not just R_i. Voters also report another signal that is used for coordination. If all voters report the same $\mathbf{R}$, then the mechanism returns $F(\mathbf{R})$. If exactly one voter reports something else, the mechanism uses the rest of the reported preferences to "punish" him, which guarantees that reporting the truth is a Nash equilibrium. If there is some greater disagreement, the mechanism uses the coordination signal in a clever way to make sure that at least one voter gains by deviating, so there are no other Nash equilibria.

SE implementation Let $SE_g(\mathbf{R}) \subseteq A$ be all candidates that win in some strong equilibrium of mechanism g. Recall that an equilibrium is strong if there is no subset of voters that can all gain by deviating.

Formally, a mechanism g implements a mapping $G : \mathcal{L}(A)^n \to 2^A$ in SE, if $SE_g(\mathbf{R}) = G(\mathbf{R})$ for all $\mathbf{R}$.

Note that we do not require that G is a valid SCC, as it may return the empty subset. An example of such a mapping is G_{CON}, which returns the (possibly empty) set of Condorcet winners of profile $\mathbf{R}$.

SE implementation was studied by Sertel and Sanver [72], which in particular showed the following.

Theorem 5.4. *The Plurality voting rule implements G_{CON} in SE, for all odd n.*

Proof. Suppose first that $\mathbf{R}$ has a Condorcet winner $c^* \in A$. We need to show that c^* wins in some SE, and that no other candidate wins in SE. Consider the profile where all voters vote for c^*. A subset of voters S can only change the outcome if $|S| > n/2$ voters deviate to some other candidate c'. However, since c^* is a Condorcet winner, at least $n/2$ voters strictly prefer c^* over c', and therefore at least dome of the voters in S would not gain by deviating to c'. Thus this is an SE.

On the other hand, consider any profile in which c wins and c is not a Condorcet winner. There is some c' such that $> n/2$ voters prefer c' over c. All these voters that are not currently voting to c' can deviate to make c' the winner. Thus this is not an SE. In particular, if there is no Condorcet winner in $\mathbf{R}$, then there are no SEs. $\square$

Other notions of implementation Several authors suggested other equilibrium concepts, and showed how certain voting rules can be implemented (sometimes by a simple mechanism or by the same rule) under these equilibria. Some examples are: Veto implements itself under *Protective equilibrium* [4]; Plurality implements itself under *Demanding equilibrium* [45]; and Plurality implements Maximin under maximal *Scoring rules* [22].

5.2 Truth bias

As we saw, Nash equilibrium (even pure) is not very informative in most common voting rules, as voters typically cannot change the outcome. Suppose that a voter who is non-pivotal, prefers to vote according to her true preferences. Formally, we can think of truth-bias as a negligible non-zero cost attached to any vote that is non-truthful (say, the mental burden of finding a strategy, or the moral burden of manipulation).

Hopefully, even such a small behavioral bias may help eliminating many undesired equilibria (e.g., when all voters vote for bottom candidates). Truth-bias was formalized independently in [43] and [20], who also provided some examples. In general, truth bias may eliminate *all* Nash equilibria. Consider for example Plurality with lexicographic tie-breaking, where voters preferences are $c \succ_1 a \succ_1 b \succ_1 d$ and $d \succ_2 b \succ_2 a \succ_2 c$. One can check that none of the 16 action profiles are stable under truth-bias.

Obraztsova et al. [58] provide characterizations of Nash equilibria in Plurality with truth bias. In particular they show:

Theorem 5.5 (Obraztsova et al. [58])**.** *Given a preference profile* $\mathbf{R}$*, and* $c = f(\mathbf{R})$ *the truthful outcome of Plurality:*

1. *It can be tested efficiently if* $\mathbf{R}$ *is an equilibrium under truth bias.*
2. *Given* $c' \neq c$*, it is NP-hard to decide if there is some action profile* $\mathbf{R}'$ *s.t.* $f(\mathbf{R}') = c'$ *and* $\mathbf{R}'$ *is an equilibrium under truth bias.*

Thompson et al. [75] studied the effect of truth bias using simulations. They generated many preference profiles and computed *all* their Nash equilibria (for Plurality) with and without truth bias. First of all, they showed that instances where no PNE exist are uncommon, and especially rare when there is an even number of voters. Yet, the number of PNEs drops dramatically from millions (i.e. almost all action profiles) to typically under 5. In addition they show that the outcome in these equilibria is more socially efficient according to various criteria.

We should note that as the number of voters grow, the probability that the truthful profile is a PNE tends to 1, both with and without truth-bias.

Lazyness and the no-vote paradox Suppose that instead of experiencing ε cost for voting strategically, voters experience ε cost for voting at all, and may choose to abstain (or stay home). Since in most profiles all or almost all voters are non-pivotal, introducing such "lazy-bias" means that almost no voter has any incentive to show up, an observation widely known as the *no-vote paradox* [59, 51].

5.3 Equilibrium under uncertainty

We have seen that when voters know exactly how others are going to vote, they rarely have an incentive to strategize (or to vote at all). Yet we known that people often do vote strategically, or at least trying to. One possible explanation for this discrepancy is *uncertainty*: since voters do not know exactly the preferences and actions of others, they know they *might* be pivotal, and hence some actions may be better than others in expectation.

There are many models introducing uncertainty to voting and we will only cover several prominent ones.

Bayesian reasoning The classic game-theoretic approach for uncertainty, or *games with partial information*, assumes that each player's *type* (preferences) is sampled from some distribution and this distribution is common knowledge. Thus each player knows her own type, and some distribution on the other players' types. In (Bayes-Nash) equilibrium, each player is playing a mixed strategy contingent on her type, that is a best response to the (mixed) joined strategy of all other players.

Myerson and Weber adapted this line of reasoning for strategic voting [52]. In their model, voters' types are cardinal utility scales for all candidates: Let $\mathcal{D}$ be a distribution over $\mathbb{R}^m$, then for each voter u is sampled i.i.d. from $\mathcal{D}$. The crucial observation is that for most common voting rules, and in particular for PSRs, there is no need to guess the exact actions of all other voters. To assess the value of every action, it is sufficient to consider the probability that the voter is pivotal between every pair of candidates. In scoring-based voting rules, each rule can be defined by specifying a set of valid scores $V \subseteq \mathbb{R}^m_+$.[6]

Formally, suppose that for any $x, y \in A$, p_{xy} is the probability that x, y are tied, and all other candidates have a strictly lower score. Myerson and Weber make a simplifying assumption here that if the voter assigns $v_y > v_x$ points to candidates y, x respectively, then the probability the winner changes from x to y is proportional to $p_{xy}(v_y - v_x)$. The expected gain for the voter from voting v when her preferences are u (utility compared to abstaining):

$$gain(\mathbf{p}, u, v) = \sum_{x,y \in C} [\text{probability outcome changes from } x \text{ to } y](u_y - u_x) = \sum_{x,y \in C} p_{xy}(v_y - v_x)(u_y - u_x).$$

In Plurality this can be somewhat simplified. The gain from voting for y is:

$$gain(\mathbf{p}, u, y) = \sum_{x \in C} [\text{probability outcome changes from } x \text{ to } y](u_y - u_x) = \sum_{x \in C} p_{xy}(u_y - u_x).$$

From here, Myerson and Weber adopt the standard game-theoretic assumptions: each voter is voting in a way that maximizes her expected gain. Thus an *equilibrium* can be described by an action profile $\mathbf{v} : N \to V$, candidates' scores $\mathbf{s} \in \mathbb{R}^m_+$, and a probability vector $\mathbf{p} \in \Delta(A^2)$, such that:

- For every voter $i \in N$ with utility u_i, $v_i = \text{argmax}_{v \in V}\, gain(\mathbf{p}, u, v)$;
- For every candidate $x \in C$, $s(x) = \sum_{i \in N} v_i(x)$;
- $p_{xy} \sim \frac{1}{k}$ if (x, y) is one of the pairs whose probability to be tied is maximal (k is the number of such pairs), and otherwise p_{xy} is (close to) 0.

There are two cases determining which ties are most likely. Let $W \subseteq C$ be such that $s(w) = s$ for all $w \in W$ and $s(c) < s$ for all other candidates (the winner set). Case (a): $|W| > 1$, in which case the most likely ties are among all $k = \binom{|W|}{2}$ pairs from W; (b) $W = \{w\}$, in which case the most likely ties are between w and any of the k candidates whose score is maximal among $C \setminus W$.

Intuitively, $\mathcal{D}$ is highly concentrated around a particular utility profile $\mathbf{u}$. As the model assumes the number of voters tends to infinity, all ties that are not the most probable ones according to $\mathbf{s}$ can be neglected. A natural question is whether such an equilibrium exists.

Theorem 5.6 (Myerson and Weber [52])**.** *For any scoring-based rule (including SCRs), and utility profile* $\mathbf{u}$*, there exists a voting equilibrium.*

[6] This captures all PSRs, Approval, Range voting and so on. E.g. in Plurality V contains vectors with a single '1' and all other entries are '0', whereas in Range voting $V \subseteq \mathbb{R}^m_+$.

In fact, for Plurality the existence of equilibrium is trivial: consider any two candidates $x, y \in C$, then there is an equilibrium where each voter votes for her more preferred candidate among x, y.

Trembling hand perfection Consider a simpler model with a finite number of voters, and the Plurality rule. Even if the voting profile is known, we can introduce uncertainty as follows. For every voter $i \in N$, we do not count the vote of i with some small independent probability ε. This may be justified as a probability that the vote is 'miscounted' or that the voter fails to vote [46].

For an action profile $\mathbf{a} \in A^n$, denote by $\mathbf{p}_\varepsilon(\mathbf{a})$ the distribution over profiles attained by eliminating each vote with i.i.d probability ε. Note for example that voting for i's least preferred candidate b is never a best response, since with positive probability b is tied with some other candidate, in which case i strictly loses, whereas i never strictly gains from voting to b.

Thus for a preference profile $\mathbf{R}$, a *robust equilibrium* (RE) is an action profile $\mathbf{a} \in A^n$ s.t. each a_i is a best-response to $\mathbf{p}_\varepsilon(\mathbf{a})$ for $\varepsilon \to 0$. This notion is closely related to *trembling hand equilibrium* in the game theory literature [65].

Messner and Polborn show that in any Robust equilibrium under Plurality, at most 2 candidates receive a positive number of votes. This phenomenon is known as the Duverger Law [69].

In fact. it is not hard to see that for any two candidates $x, y \in C$, there is an RE where only x and y get votes: this is the profile where each voter votes for the more preferable candidate among x, y.

A similar model of trembling-hand equilibrium was more recently analyzed in [55]. The main difference is that with probability ε a voter submits a random ballot, rather than fails to vote at all.

6 Iterative Voting

In the iterative voting model, voters have fixed preferences and start from some announcement (e.g., sincerely report their preferences). Votes are aggregated via some predefined rule (e.g. Plurality), but can change their votes after observing the current announcements and outcome. The game proceeds in turns, where a single voter changes his vote at each turn, until no voter has objections and the final outcome is announced. This process is similar to online polls via Doodle or Facebook, where users can log-in at any time and change their vote. Similarly, in offline committees the participants can sometimes ask to change their vote, seeing the current outcome.

The outcome of such a game depends on the exact specification of the voting rule and the process, and also on the strategies that voters use. If we assume voters do not know others' preferences or who might change their vote, then a plausible assumption is that voters will act in a myopic way. That is, vote in every round as if it is the last one. In game-theoretic terms, each voter will play a best-reply to the *current action profile* of the other voters. If no voter wants to change her vote, then by definition the current profile is a PNE. This is the type of behavior we will consider in this section. Voters who are more sophisticated on one hand, or have less accurate information about the current state on the other hand, may not follow their best-reply and instead use other heuristics. We consider such heuristics in the next section.

For either voting rule and type of behavior we are interested in the following questions:

- Are voters guaranteed to converge, i.e. reach a stable state where no voter wants to move?
- If so, how fast?
- Can we characterize the stable states?
- Is the iterative process leading the society to a socially good outcome?

6.1 Convergence and acyclicity

Local improvement graphs and schedulers Any game G induces a directed graph whose vertices are all action profiles (states) $\mathcal{A}$, and edges are all local improvement steps [78, 2]. That is, there is an edge from profile $\mathbf{a}$ to profile $\mathbf{a}'$ if there is some agent i and some action a'_i that $\mathbf{a}' = (\mathbf{a}_{-i}, a'_i)$ and i prefers $\mathbf{a}'$ to $\mathbf{a}$.

The *sinks* of G are all states with no outgoing edges. Clearly, a state is a sink iff it is a PNE. Since a state may have multiple outgoing edges ($|I(\mathbf{a})| > 1$), we need to specify which one is selected in a given play.

A *scheduler* ϕ selects which edge is followed at state $\mathbf{a}$ at any step of the game [3]. The scheduler can be decomposed into two parts, namely selecting an agent i to play (agent scheduler ϕ^N), and selecting an action in $I_i(\mathbf{a})$ (action scheduler ϕ^A), where $\phi = (\phi^N, \phi^A)$. We note that a scheduler may or may not depend on the history or other factors, but this does not affect any of our results.

Acyclicity properties Given a game G, an initial action profile $\mathbf{a}^0$ and a scheduler ϕ, we get a unique (possibly infinite) path of steps.[7] Also, it is immediate to see that the path is finite if and only if it reaches a Nash equilibrium (which is the last state in the path). We say that the triple $\langle G, \mathbf{a}^0, \phi \rangle$ *converges* if the induced path is finite.

Much attention has been given in the game theory literature to the question of convergence, and several notions of convergence have been defined [48, 47, 32, 3].

Definition 6.1. *A game G has the* finite individual improvement property *(we say that G is FIP), if $\langle G, \mathbf{a}^0, \phi \rangle$ converges for* any $\mathbf{a}^0$ *and scheduler ϕ. Games that are FIP are also known as* acyclic games *and as* generalized ordinal potential games *[48]. Two weaker notions of acyclicity are as follows.*

- *A game G is* weakly-FIP *if there is* some *scheduler ϕ such that $\langle G, \mathbf{a}^0, \phi \rangle$ converges for any $\mathbf{a}^0$. Such games are known as* weakly acyclic*, or as ϕ-potential games.*
- *A game G is* restricted-FIP *if there is* some action scheduler *ϕ^A such that $\langle G, \mathbf{a}^0, (\phi^N, \phi^A) \rangle$ converges for any $\mathbf{a}^0$ and ϕ^N [32]. We term such games as* order-free acyclic.

Intuitively, restricted FIP means that there is some restriction players can adopt s.t. convergence is guaranteed regardless of the order in which they play. Kukushkin identifies a particular restriction of interest, namely restriction to best-reply improvements, and defines the *finite best-reply property* (FBRP) and its weak and restricted analogs. We emphasize that an action scheduler *must* select an action in $I_i(\mathbf{a})$, if one exists. Thus restricted dynamics that may disallow all available actions do not fall under the definition of restricted-FIP.

[7] By "step" we mean an individual improvement step, unless specified otherwise.

$f^{PL}_{\mathbf{w},\hat{\mathbf{s}}}$	a	b	c
a	$(14,9,3)$ $\{a\}$	$(10,13,3)$ $\{b\}$	$(10,9,7)$ $\{a\}$
b	$(11,12,3)$ $\{b\}$	$(7,16,3)$ $\{b\}$	$(7,12,7)$ $\{b\}$
c	$(11,9,6)$ $\{a\}$	$(7,13,6)$ $\{b\}$	$(7,9,10)$ $\{c\}$

Figure 3: A game form $f^{PL}_{\mathbf{w},\hat{\mathbf{s}}}$, where $N = \{1,2\}$, $A_1 = A_2 = C = \{a,b,c\}$, $\hat{\mathbf{s}} = (7,9,3)$ and $\mathbf{w} = (3,4)$ (i.e., voter 1 has weight 3 and voter 2 has weight 4). The table shows the final score vector $\mathbf{s}_{(a_1,a_2)}$ for every joint action of the two voters, and the respective winning candidate $f^{PL}_{\mathbf{w},\hat{\mathbf{s}}}(a_1,a_2)$ in curly brackets.

$\langle f, \mathbf{Q}^1 \rangle$	a	b	* c
* a	$\{\mathbf{a}\}$ $\mathbf{3,2}$	$\{b\}$ $2,1$	* $\{\mathbf{a}\}$ $\mathbf{3,2}$
b	$\{b\}$ $2,1$	$\{\mathbf{b}\}$ $\mathbf{2,1}$	$\{b\}$ $2,1$
c	$\{a\}$ $3,2$	$\{b\}$ $2,1$	$\{c\}$ $1,3$

Figure 4: A game $G = \langle f, \mathbf{Q}^1 \rangle$, where $f = f^{PL}_{\mathbf{w},\hat{\mathbf{s}}}$ is as in Fig. 3, and $\mathbf{Q}^1$ is defined by $a$ord1$b$ord1c and $c$ord2$a$ord2b. The table shows the ordinal utility of the outcome to each agent, where 3 means the best candidate. **Bold** outcomes are the NE points. Here the truthful vote (marked with *) is also a NE.

We say that a game G is *FIP from state* $\mathbf{a}$ if $\langle G, \mathbf{a}, \phi \rangle$ converges for any ϕ. Clearly a game is FIP iff it is FIP from $\mathbf{a}$ for any $\mathbf{a} \in A^n$. The definitions for other all other notions of finite improvement properties are analogous.

Recall that every voting rule (game form) f and preference profile $\mathbf{R}$ induce a game. We say that f is FIP if for *any* preference profile $\mathbf{R}$ the induced game $(f, \mathbf{R})$ is FIP, and similarly for the other convergence notions.

6.2 Plurality

We allow for a broader set of "Plurality game forms" by considering both weighted and fixed voters. Each of the strategic voters $i \in N$ has an integer weight $w_i \in \mathbb{N}$. In addition, there are $\hat{n}$ "fixed voters" who do not play strategically or change their vote. The vector $\hat{\mathbf{s}} \in \mathbb{N}^m$ (called "initial score vector") specifies the number of fixed votes for each candidate. Weights and initial scores are part of the game form.[8]

The *final score* of c for a given profile $\mathbf{a} \in C^n$ in the Plurality game form $f_{\mathbf{w},\hat{\mathbf{s}}}$ is the total weight of voters that vote c. We denote the final score vector by $\mathbf{s}_{\hat{\mathbf{s}},\mathbf{w},\mathbf{a}}$ (often just $\mathbf{s}_{\mathbf{a}}$ or $\mathbf{s}$ when the other parameters are clear from the context), where $s(c) = \hat{s}(c) + \sum_{i \in N: a_i = c} w_i$.

Thus the Plurality rule selects some candidate from $W = \operatorname{argmax}_{c \in C} s_{\hat{\mathbf{s}},\mathbf{w},\mathbf{a}}(c)$ with the lowest lexicographic index. As with $\mathbf{s}$, we omit the scripts $\mathbf{w}$ and $\hat{\mathbf{s}}$ when they are clear from the context.

[8] This assumption only makes results more general, see [42].

$\langle f, \mathbf{Q}^2 \rangle$	a	b	$*\ c$
$*\ a$	$\{a\}\ 3,1$	$\{\mathbf{b}\}\ \mathbf{1,2}$	$*\ \{a\}\ 3,1$
b	$\{b\}\ 1,2$	$\{\mathbf{b}\}\ \mathbf{1,2}$	$\{b\}\ 1,2$
c	$\{a\}\ 3,1$	$\{b\}\ 1,2$	$\{c\}\ 2,3$

Figure 5: This game has the same game form as in Fig. 3, and the preference profile $\mathbf{Q}^2$ is $a$ord1$c$ord1b and $c$ord2$b$ord2a. In this case, the truthful vote $\mathbf{a}^*(\mathbf{Q}^2)$ is not a NE.

Unfortunately, Plurality is not acyclic, and this holds even if voters are restricted to best-replies.

Proposition 6.2. *f^{PL} is not FBRP.*

Proof. $C = \{a, b, c\}$. We have a single fixed voter voting for a, thus $\hat{\mathbf{s}} = (1, 0, 0)$. The preference profile is defined as $a$ord1$b$ord1c, $c$ord2$b$ord2a. The following cycle consists of better replies (the vector denotes the votes (a_1, a_2) at time t, the winner appears in curly brackets):

$$(b, c)\{a\} \xrightarrow{2} (b, b)\{b\} \xrightarrow{1} (c, b)\{a\} \xrightarrow{2} (c, c)\{c\} \xrightarrow{1} (b, c).$$

□

Direct replies We identify a different restriction, namely *direct reply*, that is well defined under the Plurality rule. Formally, a step $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ is a direct reply if $f(\mathbf{a}') = a'_i$, i.e., if i votes for the new winner. An example of an indirect step is when a voter who votes for the winner changes the outcome by moving to a candidate with a low score (e.g. the steps of voter 1 in the example above).

Theorem 6.3 ([43, 39])**.** *$f^{PL}_{\hat{\mathbf{s}}}$ is FDRP. Moreover, any path of direct replies will converge after at most m^2n^2 steps. In particular, Plurality is order-free acyclic.*

6.3 Other rules

For other voting rules, results are not as rosy (at least for strong/order-free acyclicity). While Veto was shown to converge under direct replies from any initial state [34, 67], for most common voting rules it is possible to construct examples of cycles, even when voters start by voting truthfully. These results are summarized in Table 1.

Further, even variations of the Plurality rule, such as adding voters' weights and/or changing the tie-breaking method may result in games with cycles. At least for Plurality with random tie-breaking rule, it can be shown that it is *weakly acyclic*, thereby providing partial explanation to the fact that simulations almost never hit a cycle [39]. Whether other common voting rules are also weakly acyclic is an open question, which is particularly of interest for a large number of voters.

The structure of equilibria attained under iterative voting was studied in [64], which considered both the model above, and variations where voters are truth-biased or lazy.

Voting rule	FIP	FBRP	restricted-FIP	Weak-FIP
Dictator	V	V	V	V
Plurality	X	X [MPRJ10]	V [MPRJ10,M16]	V
Veto	X	X [M16]	V [RW12,LR12]	V
k-approval ($k \geq 2$)	X	X [LR12,L15]	X	X [M16]
Borda	X	X [RW12,LR12]	X	X [RW12]
PSRs (except k-approval)	X	X [LR12,L15]	?	?
Approval	X	X [M16]	V [M16]	V
Other common rules	X	X [KLR16]	?	?

Table 1: Positive results carry to the right side, negative to the left side. All rules in the table use lexicographic tie-breaking. Reference codes: MPRJ10 [43], RW12 [67], LR12 [34], M15 [38], L15 [33], M16 [39], KLR16 [31].

Other notions of convergence The above model only considers voters who change their vote one-by-one. Other iterative models exist, that make different assumptions. For example, we can consider voters that make coordinated coalitional moves [32, 28], simultaneous (non-coordinated) moves [38], or a different dynamic where in each step a voter proposes one alternative to replace the current winner using a Majority vote [1].

6.4 Welfare implications

In those cases where an iterative voting game converges, we would like to know "how good" the outcome is to the society. As with other game-theoretic analyses of voting outcomes, there are at least two different approaches to measure outcome quality:

- with respect to the particular voting rule in question.
- with respect to an objective measure, such as social welfare, Condorcet efficiency etc.

Dynamic Price of Anarchy A common way to measure the inefficiency in a game due to strategic behavior is the *Price of Anarchy*: the ratio between the quality of the outcome in the worst Nash equilibrium, and the optimal outcome [12]. In the context of voting, this translates to the question of how far the equilibrium outcome can be from the truthful voting outcome (seeing the truthful outcome as "optimal" according to the voting rule in use).[9]

As we have seen, Nash equilibria in most voting rules can be arbitrarily far from the truth, thus Branzei et al. [11] suggested instead to restrict attention to the set of Nash equilibria that are the outcome of some iterative voting procedure starting from the truthful vote.

Formally, for a given score-based voting rule f where the candidate c with the highest score $s_f(c, \mathbf{R})$ wins in action profile $\mathbf{R}$. The dynamic Price of Anarchy is defined as

$$DPoA(f) = \min_{\mathbf{R}} \min_{\mathbf{R}' \in EQ^T(f,\mathbf{R})} \frac{s_f(f(\mathbf{R}'), \mathbf{R})}{s_f(f(\mathbf{R}), \mathbf{R})},$$

where $EQ^T(f, \mathbf{R})$ is the set of all profiles $\mathbf{R}'$ s.t.:

[9] This notion of approximation is similar to the one we saw in Section 4.3, but instead of comparing the outcome under two different rules, we consider the same rule under truthful and rational behavior.

- There is a path of best-replies from the truthful profile $\mathbf{R}$ to $\mathbf{R}'$.
- $\mathbf{R}'$ is a Nash equilibrium of $(f, \mathbf{R})$.

Branzei et al. [11] show that the DPoA of Plurality is close to 1 (i.e. a winner in equilibrium must have a very close score to the truthful winner); and that the DPoA in Veto depends on the number of candidates m. In particular for $m \leq 3$ the DPoA in Veto is constant, regardless of n. The DPoA in Borda on the other hand, is $\Omega(n)$, meaning that equilibria can be arbitrarily bad.

Objective quality metrics The fact that we chose to use a particular voting rule does not necessarily mean that this rule represents the optimal outcome for every profile. The selection of the rule might be affected by other properties of the rule such as simplicity, from tradition, and so on. We may thus have multiple criteria for a "good outcome," and ask how well a given voting rule satisfies them. Note that we focus on the equilibria outcomes rather than on the truthful outcome.

For example, we may be interested in the social welfare of the voters (as measured by Borda score), in the likelihood of finding the Condorcet winner when one exists, or avoiding the Condorcet loser, and so on.

This question was studied using extensive simulations in [40] for the Plurality rule, where it was shown that equilibria outcomes are *better* than the truthful outcome under most metrics observed. [31] performed similar simulations for several other voting rules, and obtained mixed results w.r.t. the the social welfare. More interestingly, rules that are not Condorcet consistent (Bucklin, STV) are more likely to find the Condorcet winner under rational play than under truthful voting, where as the Condorcet efficiency of Condorcet-consistent rules only slightly declines.

That said, when the number of voters is large, the initial outcome is almost always an equilibrium (whether truthful or not), and thus the effect of best-reply dynamics becomes negligible.

7 Voting Heuristics

The equilibrium analysis and the best-reply dynamics considered in the previous sections makes the following implicit assumptions.

- Voters know exactly how other voters currently vote.
- Voters are myopic: always vote as if the game ends after the current turn.
- Voters are rational: always vote in a way that improves or maximizes their utility.

7.1 Ad-hoc heuristics

Most heuristics are similar to best reply in that they only assume the voters knows her own preferences, and has some information about the current voting profile (e.g., the score of each candidate, or their current ranking). However in contrast to best- or better-reply, a heuristic step may or may not change the outcome. It thus reflects the belief of the voter that she might be pivotal even if this is not apparent from the current state. In the next subsection we will look more closely on such a rational (or bounded-rational) justification, but for now we will be satisfied with just describing some heuristics that have been proposed.

In the following description, we assume f is some score-based rule, unless specified otherwise. Let R_i be the real preferences of voter i, and $\mathbf{a} = (a_1, \ldots, a_n), \mathbf{s} = (s_1, \ldots, s_m)$ be the current action profile and current scores of all candidates. We denote by c_j the candidate with the j'th highest score in $\mathbf{s}$.

Crucially, some of these heuristics depend on some internal private parameters, which can be used to explain behavioral differences among voters.

- "k-pragmatist" [66]. Here, each voter has a parameter k_i, and votes as if only candidates $W = \{c_1, \ldots, c_k\}$ participate (ranking all other candidates below according to R_i).[10]

- "Threshold". Similar to k-pragmatist, except instead of a fixed parameter k_i, the set of "possible winners" W consists of all candidates whose score s_j is above some threshold $T_i(\mathbf{a})$ (i.e., the threshold may depend both on i and on the current state).

- "Second chance" [29]: If the current winner is not i's best or second-best choice according to R_i, she moves her second-best alternative to the top position.

- "Best upgrade" [29]: This is a restriction of best-reply to candidates that are ranked above the current winner $f(\mathbf{a})$ in R_i.

- "Upgrade" [57]: Similar to Best Upgrade, except the upgraded candidate is not necessarily placed first (only high enough to win).

- "Unit upgrade" [57]: Similar to Upgrade, except the upgraded candidate is moved exactly one step up (if this is enough to win).

Simulations show that these heuristics almost always lead to convergence when applied in an iterative voting setting [29].

7.2 Strict uncertainty and bounded rationality

A different approach to derive heuristic voting behavior, is to consider a formal way to model voter's *uncertainty* regarding the outcome. Then , based on her beliefs, the voter selects the action (ballot) that is best for her.

To see why this may differ from a purely rational behavior, note that:

1. The beliefs of the voter may not be justified. This is in contrast with e.g. the Myerson and Weber model, where the distribution voters believe is itself part of the equilibrium.

2. The response of the voter may not maximize her expected utility, which may not even be well defined.

[10] There is another variation where the most preferred candidate in W is moved to the top ,without other changes.

Thus we can think of such approaches as models of voters with *bounded rationality*. In contrast with the models we considered in Section 5.3, voters may not assign exact probabilities to outcomes, and in particular cannot compute expected utilities.[11]

Several papers studied beliefs based on strict uncertainties. One that we already mentioned is [66], which assumes the voter may only be aware of the current winner, the order of candidates according to scores, etc. Similar models are are considered in [16, 76].[12] We can also think of voters whose beliefs depend on some internal parameter, which we can think of as their *uncertainty level*. Intuitively, as the voter is more uncertain, she considers more outcomes as possible.

Two such models based on voter's optimism were suggested in [68, 56]. In both models voters with greater uncertainty will consider a larger set of candidates as possible winners, and will vote strategically to one of them.

Local Dominance A third model to strict uncertainty is based on local dominance [40], and can be applied to scoring-based rules. This model explicitly separates the beliefs of the voter on candidates scores and her strategic actions:

- All voters share some prospective score vector $\mathbf{s} = (s_1, \ldots, s_n)$.
- Each voter i has an uncertainty parameter r_i.
- Voter i considers as *possible* all outcomes $\mathbf{s}'$ such that $|s_j - s'_j| \leq r_i$.[13] Denote all possible states in voting profile $\mathbf{a}$ by $S_i(\mathbf{s_a})$, where $\mathbf{a_s}$ is the score vector induced by action profile $\mathbf{a}$.
- Given this belief, voter i will change her action from a_i to a'_i if action a'_i *dominates* action a_i. Formally, if $f(\mathbf{s}', a'_i) \succeq_i f(\mathbf{s}', a_i)$ for all states $\mathbf{s}' \in S_i(\mathbf{s_a})$, and $f(\mathbf{s}'', a'_i) \succ_i f(\mathbf{s}'', a_i)$ for at least one state $\mathbf{s}'' \in S_i(\mathbf{s_a})$.

This behavior encodes bounded rationality under *loss aversion*: the voter will make a strategic move only if certain (according to her beliefs) that this move will not hurt her, and might be beneficial for her.

Denote by $W_i \subseteq C$ the set of candidates whose Plurality score (without voter i) is at least $max_c s(c) - 2r_i$. Note that these are exactly the candidates considered as possible winners by voter i, since there is a possible state $\mathbf{s}'$ where $j \in W_i$ gets r_i more votes, and the current winner gets r_i votes less.

Worst-Case Regret minimization (WCR) [38] and Non-Myopic voting (NM) [56] are similar in the way they derive the set of possible winners W_i, but then make some different behavioral assumptions on action selection, which we will not specify explicitly here.

It turns out that the local dominance provides a (bounded) rational justification to the threshold heuristic we described above, at least in Plurality.

Lemma 7.1 ([40, 38])**.** *a'_i locally dominates a_i only if:*

[11] Also note that expected utility is undefined for a voter with ordinal preferences, even if we had such a distribution.

[12] In all these models, voters' strategy is based on dominance (similarly to the Local Dominance model presented next), but the the structure of voters' beliefs is substantially different.

[13] The paper also considers other distance metrics, but this is the simplest one.

1. *Either $a_i \notin W_i$, or a_i is the least preferred candidate in W_i;*

2. *$a'_i \in W_i$;*

3. *there is some $c \in W_i$ that is less preferred than a'_i.*

In addition, if the above conditions apply, then the most preferred candidate $a'_i \in W_i$ always locally dominates a_i.

Therefore, a voter that simply follows the threshold heuristics is essentially strategizing according to local dominance.[14] If the voter selects a step *minimizing her worst-case regret* rather than following local dominance moves, then this coincides with the threshold heuristics exactly [38].

Local dominance is strongly related to models based on modal logic and epistemology [70, 76].

7.3 Equilibrium and Convergence

Given a voting rule f and a population of voters with well defined heuristics, a *voting equilibrium* is simply a profile of valid votes $\mathbf{a}$, such that for each $i \in N$, the heuristic action of i in profile $\mathbf{a}$ is her current action a_i. That is, a state where no voter wants to change her vote.

Observation 7.2. *Consider an arbitrary voting rule f, and any restricted better-reply dynamics (including best-reply and better-reply). Then for any preferences profile $\mathbf{R}$, a state $\mathbf{a}$ is a voting equilibrium if and only if $\mathbf{a}$ is a pure Nash equilibrium of the game $(f, \mathbf{R})$.*

Given a voting rule and a heuristic, two important questions are: (a) does an equilibrium exist? and (b) will voters converge to equilibrium? The latter question can be further split to whether convergence is guaranteed from arbitrary initial states or from the truthful state.

Recall that the FIP property means that convergence is guaranteed regardless of the initial state, the order of voters, and which available reply they choose. As any heuristic simply replaces the (possibly empty) set of better replies with some other set, we can modify the definition of FIP, or FIP from the truth, accordingly.

These questions were studied in several recent papers. Some heuristics are very easy to analyze. For example, when voters start from the truthful vote, then voters using the k-pragmatist or the Second Chance heuristics will move at most once [66, 29]. Therefore, FIP from the truth is immediate.

Obraztsova et al. [57] identified some common structure for heuristic dynamics, which can be used to prove convergence for various combinations of voting rules and heuristics. However, all these studies are restricted to voters that start by reporting the truth, and use exactly the same heuristics. Results are summarized in Table 2.

[14] Note that while any LD move is consistent with the threshold heuristics, the converse does not always hold. E.g. if there 5 possible winners above the threshold, then a move from the third-preferred to the most preferred is not a local dominance move. Indeed, a loss-averse voter would refrain from such a dangerous move.

Voting rule	k-pragmatist	Second Chance	Best upgrade	Upgrade	Unit Upgrade
PSRs	V [RE12]	V	V [GLR+13]	?	V* [OMM+15]
Maximin	V [RE12]	V	V [GLR+13]	V [OMM+15]	V [OMM+15]
Copland	V [RE12]	V	V [GLR+13]	?	?
Bucklin	-	V	?	?	V [OMM+15]
all rules	-	V [GLR+13]	?	?	?

Table 2: Positive results means FIP from the truth for uniform population (where relevant). Convergence for other conditions was not studied. (* common PSRs including Borda and Plurality). Reference codes: RE12 [29], GLR+13 [29], OMM+15 [57].

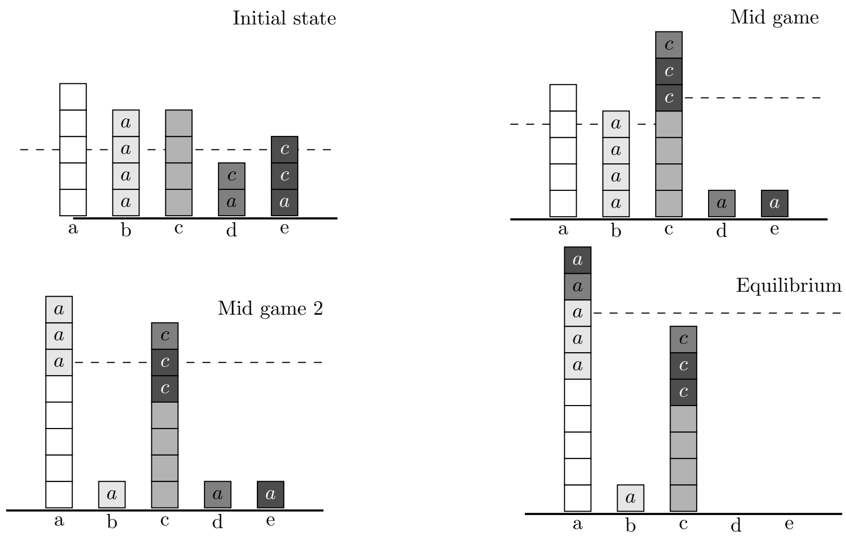

Figure 6: The top left figure shows the initial (truthful) state of the game. The letter inside a voter is his second preference. The dashed line marks the threshold of possible winners W_i for voters of type $r_i = 2$. Note that due to tie breaking it is not the same for all candidates. For example, since a beats b in tie-breaking, b needs 2 more votes to win in the initial state. In the next two figures we can see voters leaving their candidates (who are not possible winners for them) to join one of the leaders. The last figure shows an equilibrium that was reached.

Uncertainty-based heuristics Uncertainty-based heuristics are more involved, especially when the society is composed of voters with different uncertainty levels. Therefore, they have been studied mostly for the Plurality rule. On the other hand, it turns out that the Local Dominance heuristics has very strong convergence properties. An example of a Plurality game where voters use the Local Dominance heuristics is in Figure 6.

Theorem 7.3 ([40, 38])**.** *Plurality with the Local Dominance heuristics is FIP. This holds for any population of voters with diverse uncertainty levels.*

Limited convergence properties were also shown for other uncertainty-based heuristics. We summarize them in Table 2.

Heuristic	Population	FIP	FIP from truth	equilibrium exists
LD	uniform	V	V [MLR14]	V
LD	uniform + truth bias	?	V[MLR14]	V
LD	diverse	V [M15]	V	V
WCR	uniform	?	V [M15]	V
WCR	diverse	X	X	X [M15]
NM	uniform	?	V [OLP+16]	V
NM	diverse	X	X [OLP+16]	?

Table 3: Convergence results for Local-Dominance and Worst-Case Regret minimization. All results are for Plurality. MLR14 [40], M15 [38], OLP+16 [56].

7.4 Equilibrium properties

Equilibrium properties are typically studied using simulations, so that uncommon or unlikely equilibria can be ignored. Simulations are carried by generating preference profiles from some distribution (e.g. Impartial culture, Urn, Plackett-Luce, etc.), set the initial profile (truthful or other), and then sample voters randomly to make a heuristic move, until an equilibrium is reached. For some heuristics voters' parameters should also be decided up front.

It should first be noted that in practice, convergence to equilibrium is achieved in practice (i.e. in simulations) almost always, whether or not this is guaranteed by theorems. Moreover, this convergence is typically very quick.

The ad-hoc heuristics we mentioned typically lead to an better winner in terms of Condorcet efficiency and Borda score [29]. As with the best-response simulations mentioned in Section 6.4, this improvement is mild, possibly since the aforementioned heuristics give rise to a small amount of strategic behavior, and thus in many profiles the equilibrium is simply the initial state (note that some of these heuristics are just restrictions of best-response).

Extensive simulations of Plurality voting with Local Dominance heuristics show the following [40]:

- As uncertainty level r increases, there is more strategic interaction among voters (more moves) until a certain point, from which strategic interaction declines.
- With more strategic interaction, the welfare measures (including Borda score, Condorcet consistency and others) tend to improve, reaching a significant improvement around the peak of strategic activity.
- With more strategic interaction, votes are more concentrated around two candidates (Duverger Law).

These findings are consistent among a broad class of preference distributions, and for different numbers of voters and candidates. Therefore at least for Plurality, it seems that equilibria reached under Local Dominance resemble outcomes we observe in reality, and avoids unreasonable or highly inefficient Nash equilibria.

8 Summary: Towards a complete theory of strategic voting

Social choice is perhaps the oldest topic that received formal game-theoretic analysis, much before the term game theory was coined. Yet while economists, political scientists, mathematicians (and now computer scientists) all agree that Nash equilibrium is not an appropriate solution concept for voting, there does not seem to be a single acceptable theory for strategic voting.

This might be due to the fact that, as in other cases that concern with human behavior, strategic voting involves many factors. Some of these factors may be domain specific and/or depend on complex cognitive and social processes. One attempt to offer a desiderata for models of strategic voting appears in [40] (a longer version is in [41]). This is a list of criteria that a successful voting theory should satisfy, and is parted into: *theoretic criteria*, aiming to keep the model grounded in game theoretic concept; *behavioral criteria*, whose aim is to maintain plausible assumptions; and *scientific criteria*, that requires the theory to be predictive and reproducible.

The material covered in this course was limited to theory, with the exception of some simulations. However to better understand strategic voting, theoretical models must be confronted with, and influenced by, empirical data and behavioral experiments [24, 7, 60, 30, 37, 74].

References

[1] Stéphane Airiau and Ulle Endriss. Iterated majority voting. In *International Conference on Algorithmic Decision Theory*, pages 38–49. Springer, 2009.

[2] Daniel Andersson, Vladimir Gurvich, and Thomas Dueholm Hansen. On acyclicity of games with cycles. *Discrete Applied Mathematics*, 158(10):1049–1063, 2010.

[3] Krzysztof R. Apt and Sunil Simon. A classification of weakly acyclic games. In *SAGT'12*, pages 1–12. 2012.

[4] Salvador Barberà and Bhaskar Dutta. General, direct and self-implementation of social choice functions via protective equilibria. *Mathematical Social Sciences*, 11(2):109–127, 1986.

[5] J. Bartholdi and J. Orlin. Single Transferable Vote resists strategic voting. *Social Choice and Welfare*, 8:341–354, 1991.

[6] J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.

[7] Anna Bassi. Voting systems and strategic manipulation: an experimental study. Technical report, mimeo, 2008.

[8] Nadja Betzler, Rolf Niedermeier, and Gerhard J Woeginger. Unweighted coalitional manipulation under the borda rule is np-hard. In *IJCAI*, volume 11, pages 55–60. Citeseer, 2011.

[9] Eleanor Birrell and Rafael Pass. Approximately strategy-proof voting. Technical report, DTIC Document, 2011.

[10] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[11] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.

[12] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 67–73. ACM, 2005.

[13] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[14] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.

[15] Vincent Conitzer and Tuomas Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *AAAI*, volume 6, pages 627–634, 2006.

[16] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. *arXiv preprint arXiv:1106.5448*, 2011.

[17] Jessica Davies, George Katsirelos, Nina Narodytska, and Toby Walsh. Complexity of and algorithms for borda manipulation. In *AAAI*, volume 11, pages 657–662, 2011.

[18] Elad Dokow and Dvir Falik. Models of manipulation on aggregation of binary evaluations. *arXiv preprint arXiv:1201.6388*, 2012.

[19] Elad Dokow, Michal Feldman, Reshef Meir, and Ilan Nehama. Mechanism design on discrete lines and cycles. In *Proceedings of 13th ACM-EC*, pages 423–440, 2012.

[20] Bhaskar Dutta and Jean-François Laslier. Costless honesty in voting. In *10th International Meeting of the Society for Social Choice and Welfare, Moscow*, page 116, 2010.

[21] Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 371–380. ACM, 2009.

[22] Dvir Falik, Reshef Meir, and Moshe Tennenholtz. On coalitions and stable winners in plurality. In *WINE'12*, pages 256–269. 2012.

[23] Robin Farquharson. *Theory of Voting*. Yale Uni. Press, 1969.

[24] Robert Forsythe, Thomas Rietz, Roger Myerson, and Robert Weber. An experimental study of voting rules and polls in three candidate elections. *International Journal of Game Theory*, 25(3):355–383, 1996.

[25] Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 243–249. IEEE, 2008.

[26] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–601, 1973.

[27] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45:665–681, 1977.

[28] Laurent Gourvs, Julien Lesca, and Anaelle Wilczynski. Strategic voting in a social context: considerate equilibria. In *ECAI'16*, 2016.

[29] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Algorithmic Decision Theory*, pages 181–192. Springer, 2013.

[30] Michael Kearns, Stephen Judd, Jinsong Tan, and Jennifer Wortman. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5):1347–1352, 2009.

[31] Aaron Koolyk, Omer Lev, and Jeffrey S Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1329–1330. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[32] Nikolai S. Kukushkin. Acyclicity of improvements in finite game forms. *International Journal of Game Theory*, 40(1):147–177, 2011.

[33] Omer Lev. *Agent Modeling of Human Interaction: Stability, Dynamics and Cooperation.* PhD thesis, The Hebrew University of Jerusalem, 2015.

[34] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proceedings of 11th AAMAS*, pages 611–618, 2012.

[35] Christian List and Clemens Puppe. Judgement aggregation: A survey. In P. Pattanaik, P. Anand, and C. Puppe, editors, *The Handbook of Rational and Social Choice*. Oxford University Press, 2009.

[36] Eric Maskin. Nash equilibrium and welfare optimality. *The Review of Economic Studies*, 66(1):23–38, 1999.

[37] Nicholas Mattei, James Forshee, and Judy Goldsmith. An empirical study of voting rules and manipulation with large datasets. *Proc. of COMSOC-12*, 2012.

[38] Reshef Meir. Plurality voting under uncertainty. In *AAAI'15*, pages 2103–2109, 2015.

[39] Reshef Meir. Strong and weak acyclicity in iterative voting. In *COMSOC'16*. 2016.

[40] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *EC'14*, 2014.

[41] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. *CoRR*, abs/1404.4688, 2014.

[42] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Acyclic games and iterative voting. *CoRR*, abs/1606.05837, 2016.

[43] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nick Jennings. Convergence to equilibria of plurality voting. In *Proceedings of 24th AAAI*, pages 823–828, 2010.

[44] Reshef Meir, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for strategyproof classification. *Artificial Intelligence*, 186:123–156, 2012.

[45] V Merlin and J Naeve. Implementation of social choice correspondences via demanding equilibria. In *Logic, game theory and social choice: proceedings of the international conference LGS*, volume 99, pages 264–280, 1999.

[46] Matthias Messner and Mattias Polborn. Robust political equilibria under plurality and runoff rule. 2005.

[47] Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13(1):111–124, 1996.

[48] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[49] Elchanan Mossel and Miklós Z Rácz. A quantitative gibbard-satterthwaite theorem without neutrality. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1041–1060. ACM, 2012.

[50] H. Moulin. On strategy-proofness and single-peakedness. *Public Choice*, 35:437–455, 1980.

[51] Hervé Moulin. Condorcet's principle implies the no show paradox. *Journal of Economic Theory*, 45(1):53–64, 1988.

[52] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.

[53] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9. Cambridge University Press, 2007.

[54] Matıas Núnez and Marcus Pivato. Truth-revealing voting rules for large populations. In *COMSOC'16*. 16.

[55] S Obraztsova, Z Rabinovich, E Elkind, M Polukarov, and N Jennings. Trembling hand equilibria of plurality voting. In *IJCAI'16*, 2016.

[56] Svetlana Obraztsova, Omer Lev, Maria Polukarov, Zinovi Rabinovich, and Jeffrey S Rosenschein. Non-myopic voting dynamics: An optimistic approach. mimeo, 2016.

[57] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R Jennings. On the convergence of iterative voting: how restrictive should restricted dynamics be? In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 993–999. AAAI Press, 2015.

[58] Svetlana Obraztsova, Evangelos Markakis, and David R. M. Thompson. Plurality voting with truth-biased agents. In *Algorithmic Game Theory*, pages 26–37. Springer, 2013.

[59] Guillermo Owen and Bernard Grofman. To vote or not to vote: The paradox of nonvoting. *Public Choice*, 42(3):311–325, 1984.

[60] T.R. Palfrey. Laboratory experiments in political economy. *Annual Review of Political Science*, 12:379–388, 2009.

[61] Ariel D Procaccia. Can approximation circumvent gibbard-satterthwaite? In *AAAI*, 2010.

[62] Ariel D Procaccia and Jeffrey S Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, 2007.

[63] Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proceedings of 10th ACM-EC*, pages 177–186, 2009.

[64] Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey S Rosenschein. Analysis of equilibria in iterative voting schemes. In *AAAI*, volume 15, pages 1007–1013. Citeseer, 2015.

[65] Eric Rasmusen and Basil Blackwell. Games and information. *Cambridge, MA*, 15, 1994.

[66] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of 11th AAMAS*, pages 635–644, 2012.

[67] Reyhaneh Reyhani and Mark C Wilson. Best-reply dynamics for scoring rules. In *20th European Conference on Artificial Intelligence*. IOS Press, 2012.

[68] Reyhaneh Reyhani, Mark C Wilson, and Javad Khazaei. Coordination via polling in plurality voting games under inertia. In *Proceedings of 20th ECAI*, 2012.

[69] William H Riker. The two-party system and duverger's law: an essay on the history of political science. *American Political Science Review*, 76(04):753–766, 1982.

[70] Eric Pacuit Samir Chopra and Rohit Parikh. Knowledge-theoretic properties of strategic. In *JELIA'04*, 2004.

[71] M. Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–216, 1975.

[72] Murat R. Sertel and M. Remzi Sanver. Strong equilibrium outcomes of voting games are the generalized condorcet winners. *Social Choice and Welfare*, 22:331–347, 2004.

[73] L.-G. Svensson. The proof of the Gibbard-Satterthwaite theorem revisited. Working Paper No. 1999:1, Department of Economics, Lund University, 1999. Available from: http://www.nek.lu.se/NEKlgs/vote09.pdf.

[74] Maor Tal, Reshef Meir, and Ya'akov (Kobi) Gal. A study of human behavior in online voting. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*, pages 665–673, 2015.

[75] David R. M. Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey S. Rosenschein. Empirical analysis of plurality election equilibria. In *Proceedings of 12th AAMAS*, pages 391–398, 2013.

[76] Hans van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. Strategic voting and the logic of knowledge. In *TARK'13*, 2013.

[77] Lirong Xia. Computing the margin of victory for various voting rules. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 982–999. ACM, 2012.

[78] H Peyton Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84, 1993.

A Proofs

A.1 The Gibbard-Satterthwaite theorem

Two useful lemmas The first lemma says that a strategy-proof voting rule's selected outcome remains constant for all changes to the preference profile such that candidates ranked below the winner before the change are also ranked below the winner after the change:

Lemma A.1. *(monotonicity) Let f be a strategy-proof voting rule, and let $f(\prec) = a$ for some preference profile $\prec$ and $a \in A$. Then $f(\prec') = a$ for all preference profiles $\prec'$ such that:*

$$\forall i \in N \ \ \forall x \in A \setminus \{a\} : x \prec_i a \Rightarrow x \prec'_i a$$

Proof. Starting from the preference profile $\prec$ we will change the voters' preferences one at a time to $\prec'$ showing that the winner remains constant at every step. We begin with the preference profile $(\prec'_1, \prec_{-1})$ and assume $f(\prec'_1, \prec_{-1}) = b$. Since f is strategy-proof, we know that $b \preceq_1 a$ (otherwise voter 1 would benefit from reporting his preference as $\prec'_1$ instead of $\prec_1$), so according to the lemma premise we get $b \preceq'_1 a$. If $b \prec'_1 a$ then voter 1 would benefit from reporting his preference as $\prec_1$ instead of $\prec'_1$, contradicting the assumption that f is strategy-proof. Therefore $a = b$, meaning that the winner has not changed due to voter 1's change in preference. In the same manner we show that the winner remains constant as each voter $i \in N$ changes his preference from $\prec_i$ to $\prec'_i$. We conclude that $f(\prec') = a$. □

The second lemma says that the outcome of a strategy-proof and onto voting rule must be (weakly) Pareto optimal, meaning there is no candidate strictly preferred by all voters to the winning candidate:

Lemma A.2. *(Pareto optimality) Let f be a strategy-proof voting rule which is onto, and let $a, b \in A, a \neq b$. If $\prec$ is a preference profile such that $\forall i \in N : b \prec_i a$ then $f(\prec) \neq b$.*

Proof. Suppose that $f(\prec) = b$. Since f is onto, there exists a preference profile $\prec'$ such that $f(\prec') = a$. Let $\prec''$ be a preference profile where all voters rank candidate a first and candidate b second. This means that no voter ranked b lower in $\prec''$ than they did in $\prec$ (based on the assumption that all voters ranked a above b in $\prec$), so by monotonicity it follows that $f(\prec'') = f(\prec) = b$. On the other hand, no voter ranked a lower in $\prec''$ than they did in $\prec'$ (since a is always ranked first in $\prec''$), so by monotonicity it follows that $f(\prec'') = f(\prec') = a$, which is a contradiction (since $a \neq b$ and f is a function). Hence $f(\prec) \neq b$. □

Theorem A.3. *If there are exactly two voters and at least three candidates, any voting rule f that is strategy-proof and onto is dictatorial.*

Proof. Let $\prec$ be a preference profile and let $a, b \in A$ such that:

$$\forall x \in A \setminus \{a, b\} : (x \prec_1 b \prec_1 a) \wedge (x \prec_2 a \prec_2 b)$$

By Pareto optimality, $f(\prec) \in \{a, b\}$. Assume WLOG that $f(\prec) = a$. Now consider a preference $\prec_2'$ which satisfies:

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

Due to Pareto optimality, $f(\prec_1, \prec_2') \in \{a, b\}$. Due to strategy-proofness, $f(\prec_1, \prec_2') \neq b$. And so $f(\prec_1, \prec_2') = a$. Monotonicity now implies that f will select a as the winner for any preference profile where voter 1 ranks a first. So voter 1 is a dictator for candidate a. The analysis above can be repeated for all pairs of candidates $x, y \in A$, to show that either voter 1 is a dictator for candidate x or voter 2 is a dictator for candidate y.

For $i \in \{1, 2\}$, let A_i denote the set of candidates for whom voter i is a dictator. Let $A_3 = A \setminus (A_1 \cup A_2)$. Note that $|A_3| \leq 1$ (otherwise we could repeat the analysis above for two candidates in A_3 and find that one of them must be in A_1 or A_2). Note that for two different candidates x, y, it is not possible that $x \in A_1$ and $y \in A_2$ (this would cause a contradiction if voter 1 ranks x first and voter 2 ranks y first). We know that $|A_3| \leq 1$ and $m \geq 3$ and therefore $|A_1 \cup A_2| \geq 2$, but we showed that two different candidates x and y cannot belong to A_1 and A_2 respectively, so $A_1 \cap A_2 = \emptyset$. It follows that $A_2 = \emptyset$ (we assumed $a \in A_1$). Finally, $A_3 = \emptyset$, since if we assume $c \in A_3$ then repeating the above analysis for c and a implies that either $c \in A_1$ or $a \in A_2$, which is a contradiction. We conclude that $A_1 = A$ and f is dictatorial with voter 1 as the dictator. (If we had initially assumed that $f(\prec) = b$, voter 2 would have been the dictator.) □

A.2 A voting rule where the manipulation is NP-Hard

Definition: Plurality with Preround (PwP)

A voting rule which consists of two rounds:

1. The candidates are divided into pairs. The loser of pairwise elections in each pair is removed.
2. After we are left with half of the original candidates, perform **Plurality**. Meaning, every voter gives one point to his top candidate who survived the 1^{st} round.

Theorem A.4. MANIPULATION$_{\mathbf{PwP}}$ $\in NP - Complete$

Proof. Obviously, $\mathbf{MANIPULATION_{PwP}} \in NP$, since given $\prec'$, we can decide in polynomial time if p was elected or not. We'll prove that is it NP-Hard using a reduction from the SAT problem[1].

Preparations

Let $\varphi = (x_1 \vee \neg x_2) \wedge (x_2 \vee \neg x_7 \vee x_8) \wedge \ldots$ be a CNF formula.

- Denote $V = \{v : v \text{ variable in the formula}\}$ a set of variables.
- Denote $L = \{v, \neg v : v \in V\}$ set of literals.
- Denote $C = \{c : c \text{ clause in the formula }\}$ set of clauses.
- Define D to be a set of dummy candidates such that $|D| = |C| + 1$.
- Define $A = \{p\} \cup L \cup C \cup D$ the set of candidates.

We'll create an instance of $\mathbf{MANIPULATION_{PwP}}$ in the following way:
Define V with the 3 following separate and covering subsets:

- Subset 1: consists of $4|C| + 2$ voters, such that p is the top candidate in their preference, and the candidates from D are last.
- Subset 2: consists of $4|C|$ voters for every $c \in C$, such that c is the top candidate in their preference, and the candidates from D are last. In total, there are $4|C|^2$ voters in this subset.
- Subset 3: consists of 4 voters for every $c \in C$, such that in their preference the literals that appear in c are ranked 1^{st}, c is ranked after them, and all the candidates from D are last. In total, there are $4|C|$ voters in this subset.

Additionally, the voters will rank the Literals in L such that for every $v \in V$, the literal v is in tie with the literal $\neg v$ (this is possible since we did not define the rankings of the majority of the voters on the literals, and there is an even number of voters).
In the preround, we'll divide the candidates to pairs such that for every $v \in V$, v is paired with $\neg v$, and every candidate in $\{p\} \cup C$ is paired with a candidate in D.

Reduction

All the candidates in $\{p\} \cup C$ survive the 1^{st} round because all the voters rank the candidates in D last (no matter what preference the manipulator provided). Additionally, since every v and his negation are in tie, the manipulator can decide for every $v \in V$ which of the two will proceed to the 2^{nd} round.

[1]See http://en.wikipedia.org/wiki/Boolean_satisfiability_problem for details

- $\Rightarrow$ Suppose that there exists a satisfying assignment to the given CNF formula. If the manipulator's preference is such that all the literals that receive true value in the assignment will proceed to the 2^{nd} round, he'll cause p to be elected: p has at least $4|C|+2$ points in the 2^{nd} round (the 1^{st} group of voters voted for him). For each $c \in C$, at least one $l \in c$ survived the 1^{st} round (because for every clause, the manipulator passed at least one literal - the one that was satisfied in the clause). Therefore, every voter of the 3^{rd} group has at least one literal ahead of c, leaving c with only the $4|C|$ points he received from the 2^{nd} group of voters.
  Additionally, every literal $l \in L$ receives points only from the 3^{rd} group of voters, leaving him with up to $4|C|$ points (the size of the 3^{rd} group).
  We've examined all the candidates that survived the 1^{st} round, and found that even if the manipulator chooses to vote to a candidate other than p, that candidate will have up to $4|C|+1$ points, not enough to surpass p and his $4|C|+2$ points. We can conclude that if there is a satisfying assignment to the given CNF formula, there is a preference by the manipulator that can cause the election of p.

- $\Leftarrow$ Suppose there is no satisfying assignment to the given formula. Because in the 1^{st} round for each variable only one of his literals proceeds to the 2^{nd} round, for every preference of the manipulator (representing an assignment) there exists a clause $c^* \in C$ such that every one of its literals were false, and therefore didn't proceed to the 2^{nd} round.
  Therefore, c^* receives $4|C|$ points from the 2^{nd} group of voters, and 4 more points from the 4 voters in the 3^{rd} group that ranked the literals of c^* 1^{st}, and c^* after them (since these literals didn't proceed to the 2^{nd} round). We conclude that the candidate c^* has at least $4|C|+4$ points while p can have at most $4|C|+3$ points (one point from the manipulator and the rest from the 1^{st} group of voters).
  We proved that if there is no satisfying assignment, p will lose no matter the manipulator's preference, as required.

In conclusion, we proved that $\textbf{MANIPULATION}_{\textbf{PwP}} \in$ NP-Hard. $\square$