Strategic Voting

Reshef Meir

I. Introduction

In a typical voting scenario, a group of voters with diverse preferences need to collectively choose one out of several alternatives. Examples include a committee that selects a candidate for a faculty position or an award, countries in an international forum voting on the adoption of a new environmental treaty, or even automated agents that vote on the preferred meeting time on behalf of their users.

As the satisfaction of each voter is determined by the selected alternative, which is in turn affected by the actions (namely, the ballots) of others, casting a vote is in fact playing a strategic game.

The study of strategic voting is an effort to utilize *game theory*, which merits to model and predict rational behavior in a wide range of economic and social interactions, to explain and even direct the strategic decisions of voters.

This review paper is a hyper-condensed version of a book on strategic voting that is forthcoming this year. [1] The main purpose of the book is to overview the main approaches to strategic voting, in a way that makes these approaches comparable across fields and disciplines. In this paper I will mention the main directions and lines of work, but almost without going into the technical details.

Our starting point will be the seminal Gibbard-Satterthwaite theorem, which states that under a set of natural requirements, one cannot hope to construct a voting rule that is immune to strategic manipulations by the voters. This mean that there will always be situations where some voters have an incentive to misreport their true preferences. From this strong negative result emerged two lines of research. One continues to shape the boundaries and limitations of truthful voting mechanisms, by relaxing some of the assumptions that lead to the G-S impossibility result. The other line forgoes the attempt to elicit truthful votes, and instead applies game theory and equilibrium analysis to understand how strategic voters would vote in existing mechanisms.

II. Basic Notations

We denote sets by upper case letters (e.g., $A = \{a_1, a_2, \ldots\}$) and vectors by bold letters (e.g., $\boldsymbol{a} = (a_1, a_2, \ldots)$).

For a finite set X, we denote by $\mathcal{L}(X)$ the set of all linear (strict) orders over X.

a) Social choice: A voting scenario is defined by a set of *candidates*, or *alternatives*, A, a set of voters N, and a *preference profile* $\boldsymbol{L} = (L_1, \ldots, L_n)$, where each $L_i \in \mathcal{L}(A)$. For $a, b \in A, i \in N$, candidate a precedes b in L_i (denoted $a \succ_i b$) if voter i prefers candidate a over candidate b. We can also think about more general preferences, such as cardinal utilities that we denote by $U_i : A \to \mathbb{R}$.

Definition 1 (Social choice correspondence)**.** *A social choice correspondence (SCC) is a function* $F : \mathcal{L}(A)^n \to 2^A \setminus \emptyset$.

Definition 2 (Social choice function)**.** *An SCC F is* resolute *if* $|F(\boldsymbol{L})| = 1$ *for all* $\boldsymbol{L}$. *Resolute SCCs are also called* Social choice functions *(SCF). We typically denote SCFs by a lower case letter* f.

We will reserve the term *voting rule* for a SCF (i.e., a rule with a single winner) unless stated otherwise.

Some common voting rules that are mentioned in the paper are based on computing some score $s(c, \boldsymbol{L})$ for every candidate $c \in A$, and selecting the candidate with the highest score (employing some tie breaking policy if needed). For example, in *Plurality* $s(c, \boldsymbol{L})$ is the number of voters who ranked c in the first place. More generally, a *positional scoring rule* (PSR) sets $s(c, \boldsymbol{L}) = \sum_{i \in N} \alpha_{L_i^{-1}(c)}$ where α is some non-decreasing vector. The *Borda* rule is an example of a PSR where $\alpha = (m-1, m-2, \ldots, 2, 1)$.

b) Game theory:

Definition 3 (Game)**.** *A (finite, n-person, non-cooperative)* game *is a tuple* $\langle N, \mathcal{A}, \boldsymbol{u} \rangle$, *where:*

- N *is a finite set of n players, indexed by i;*
- $\mathcal{A} = A_1 \times \cdots \times A_n$, *where A_i is a finite set of actions available to player i. Each vector* $\boldsymbol{a} = (a_1, \ldots, a_n) \in \mathcal{A}$ *is called an* action profile*;*
- $\boldsymbol{u} = (u_1, \ldots, u_n)$ *where* $u_i : \mathcal{A} \to \mathbb{R}$ *is a real-valued utility (or payoff) function for player i.*

A *game form* is similar to a game, except the utilities remain unspecified. Rather, for each combination of actions, we have an abstract "outcome" from some set A. Any game form $g : \mathcal{A} \to A$ together with cardinal utility functions for each player i, induce a unique game denoted by $\langle g, \boldsymbol{U} \rangle$. This is simply the normal form game $\langle N, \mathcal{A}, \boldsymbol{u} \rangle$, where $u_i(\boldsymbol{a}) \triangleq U_i(g(\boldsymbol{a}))$ for all $i \in N$ and $\boldsymbol{a} \in \mathcal{A}$. We can similarly combine g with an ordinal preference profile $\boldsymbol{L}$ to get an ordinal game.

For example, consider the game form on Fig. 1 (right). The set of players is $N = \{1, 2\}$, where 1 selects a row and 2 selects a column; $A_1 = A_2 = \{C, D\}$ (which stand for the actions "Cooperate" and "Defect"). The game on the left (the famous *prisoner's dilemma*) is obtained by setting a cardinal utility of $U_1(a) = 3, U_2(a) = 3, U_1(b) = 0$ and so on.

A (pure) *Nash equilibrium* (NE) in game $\langle N, \mathcal{A}, u \rangle$ is an action profile $\boldsymbol{a} \in \mathcal{A}$ such that $\forall i \in N \forall a'_i \in A_i$, $u_i(\boldsymbol{a}) \geq u_i(\boldsymbol{a}_{-i}, a'_i)$. That is, every player weakly preferences her current action over any other action *assuming others do*

Reshef Meir is with the faculty of Industrial Engineering and Management, Technion-Israel Institute of Technology, e-mail: (reshefm@ie.technion.ac.il).

[1] *Stategic Voting*, Reshef Meir, Synthesis Lectures on AI and Machine Learning, Morgan-Claypool, forthcoming.

	C	D
C	$3,3$	$0,4$
D	$4,0$	$1,1$

	C	D
C	a	b
D	c	d

Figure 1. On the left - one variation of the prisoner's dilemma game. On the right, a 2×2 game form.

not change their own action. For example, the profile (D, D) is the unique NE in the prisoner's dilemma above.

A profile $\boldsymbol{a}$ is *Pareto efficient* if there is no other profile that is weakly better for all agents, and strictly better for at least one. For example, in the prisoner's dilemma, all profiles *excepts* (D, D) are Pareto efficient.

c) Game Forms are Voting Rules: Note that by taking any voting rule (=game form) f, and add a specific preference profile $\boldsymbol{L}$, we get the ordinal game $\langle f, \boldsymbol{L}\rangle$ as explained above, and we can go ahead and analyze its equilibria. However, for most common voting rules, this approach is not very informative. Consider 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. This is true even if this candidate is ranked last by all voters in $\boldsymbol{L}$, since no single voter can change the outcome. In Section VI we return to the notion of equilibrium in voting, and consider refinements and variations that are more reasonable and more useful as a solution concept.

III. Strategyproofness and the Gibbard-Satterthwaite Theorem

d) Example of a manipulation: Consider an election using the Borda voting rule with the following preference profile $\boldsymbol{L}$:

L_1	L_2	L_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! This is called a *manipulation*.

A natural question is whether there are voting rules where such manipulations are impossible. That is, where a voter can never gain from lying about her preferences.

Definition 4 (Strategyproofness)**.** *A voting rule f is* strategyproof *if no single voter can ever benefit from lying about her preferences:*

$$\forall \boldsymbol{L} \in \mathcal{L}(A)^n \;\; \forall i \in N \;\; \forall L'_i \in \mathcal{L}(A)$$
$$f(L'_i, \boldsymbol{L}_{-i}) \preceq_i f(\boldsymbol{L}).$$

For example, Plurality is strategyproof when $|A| = 2$.

A voting rule f is *dictatorial* if there is an individual (the dictator) such that i's most preferred candidate is always chosen:

$$\exists i \in N \quad \forall \boldsymbol{L} \in \mathcal{L}(A)^n : f(\boldsymbol{L}) = \text{top}(L_i).$$

A voting rule f is a *duple* if there are only two possible winners.

Theorem 1 (Gibbard [22],Satterthwaite [61])**.** *A deterministic and onto voting rule is strategyproof if and only if it is 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 (G-S) theorem. Several simple proofs can be found in [66].

The G-S theorem is considered as a strong negative result: Both dictatorial rules and duples have significant shortcomings as voting rules. A dictatorship ignores the will of all voters but one, and a duple may fail to select a candidate even if there is a unanimous agreement among voters that it is best.

e) Extensions: The negative result implied by the theorem is quite robust. Several recent papers show that the number of different profiles in which there is a manipulation is relatively large (a polynomial fraction of all profiles), unless the voting rule is very close to being dictatorial [21], [42].

When also considering manipulations by *coalitions* the situation becomes even worse. For a wide class of voting rules known as "generalized scoring rules," and which contains most common voting rules, Xia and Conitzer [71] showed that "large coalitions" (with substantially more than $\sqrt{n}$ voters) can decide the identity of the winner in almost every profile. These results were later extended by Mossel et al. [41].

Another result demonstrating the robustness of the G-S theorem is by Slinko and White [65], who showed that even if we restrict manipulations by voters to be "safe" (informally, such that any number of like-minded followers will not hurt the manipulator), this does not expand the set of strategyproof voting rules.

IV. Regaining Truthfulness in Voting

We will focus on four approaches, each of which attains truthfulness by relaxing some assumption underlying the G-S theorem. Other approaches that involve monetary payments are discussed in Section V.

A. 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). Note that not every preference profile is possible under this assumption.

More formally, a preference profile $\boldsymbol{L}$ is *single peaked* w.r.t. a linear order $\mathcal{O}$ over A, if each voter has some "peak candidate" a^*_i s.t. if x is *between* a^*_i and y then i prefers x over i. See Fig. 2 for an example.

Consider a linear order $\mathcal{O}$ over alternatives A, and a preference profile $\boldsymbol{L}$ that is single-peaked on $\mathcal{O}$. The *Median* voting rule considers the peak locations of all voters, and return their median as the winner. Consider the example in Fig. 2. The median of the five numbers $\{l_1, l_2, l_3, l_4, l_5\} = \{1, 2, 4, 5, 4\}$

L_1	$A \succ E \succ C \succ D \succ B$
L_2	$E \succ A \succ C \succ D \succ B$
L_3	$D \succ B \succ C \succ E \succ A$
L_4	$B \succ D \succ C \succ E \succ A$
L_5	$D \succ C \succ E \succ B \succ A$
L_6	$D \succ C \succ B \succ A \succ E$

L_1	L_2		L_5 L_3	L_4
A	E	C	D	B
1	2	3	4	5

Figure 2. The preferences of the first five voters are single peaked w.r.t. the order $\mathcal{O} = A \succ E \succ C \succ D \succ B$. The right figure shows the position of each of the first five voters w.r.t. the order $\mathcal{O}$. For example, $l_4 = \mathcal{O}(B) = 5$ The sixth voter L_6 is not single peaked w.r.t. $\mathcal{O}$.

is 4. Thus either voter 3 or voter 5 can be the median voter. In either case, the outcome is $\text{top}(L_3) = \text{top}(L_5) = D$, which is the *median candidate*.

Theorem 2 (Black [6]). *The Median Voter rule is strategyproof.*

A natural question is what other restrictions on preferences give rise to "median-like" mechanisms that are strategyproof. This question has been studied thoroughly [43], [27], [2], [46].

B. Complexity barriers

Even though the G-S theorem states that manipulations exist under any voting rule, a voter trying to manipulate might find it difficult to know *how* to manipulate. This observation led to the idea that some voting rules might be truthful *in practice*, assuming that voters have limited computational resources.

Bartholdi, Tovey, and Trick, formalized the following computational problem, which can be applied to any voting rule f.

$\textbf{MANIPULATION}_\mathbf{f}$: given a set of candidates A, a group of voters N, a manipulator $i \in N$, a preference profile of all voters except i $\boldsymbol{L}_{-i} = (L_1, \ldots, L_{i-1}, L_{i+1}, \ldots, L_n)$, and a specific candidate $p \in A$: Answer whether the manipulator can provide a preference L_i^* such that $f(\boldsymbol{L}_{-i}, L_i^*) = p$.

Then, they asked whether there is a voting rule f such that computing the outcome $f(\boldsymbol{L})$ is easy, but the problem $\textbf{MANIPULATION}_\mathbf{f}$ is NP-hard. Note that since the number of possible reports is $m!$, a brute-force search is typically infeasible.

At least for some voting rules, it is easy to tell whether a manipulation exists or not. E.g. in Plurality it is sufficient to let i rank p at the top of L_i', followed by all other candidates in an arbitrary order. A manipulation exists if and only if $f(\boldsymbol{L}_{-i}, L_i') = p$. Thus $\textbf{MANIPULATION}_\textbf{Plurality}$ is in $\mathcal{P}$.

Theorem 3 ([4]). *There is a voting rule f such that: I) $f(\boldsymbol{L})$ can be computed in polynomial time; II) $\textbf{MANIPULATION}_\mathbf{f}$ is an NP-Complete problem.*[2]

The original proof in [4] used a variation of Copeland, and similar hardness results hold for common voting rules such as STV [3]. 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.

A recent survey of which common voting rules are hard to manipulate by individual or by a coalition appears in [8], Section 6.4.

[2] We do not formally define here what is an NP-hard problem, and refer the reader to standard textbooks (e.g., [69]) for definitions and further discussion.

C. Randomized voting rules

It is easy to see that by allowing randomization, we can find strategyproof voting 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. The following theorem by Gibbard characterizes exactly the set of randomized strategyproof voting rules.

Theorem 4 ([23]). *A (randomized) voting rule f is strategyproof, if and only if it is a lottery over duples and strategyproof unilateral rules.*

A unilateral rule is a rule that depends on the report of a single voter (e.g., a dictatorial rule). It should be noted that in order to extend the notion of manipulation and strategyproofness to randomized outcomes, Gibbard assumes that each voter has a *cardinal* utility function U_i over alternatives, and a manipulation means that a voter gains in expectation.

Some recent work used Gibbard's characterization to derive strategyproof voting mechanisms with some desired properties. For example, Procaccia [55] proved that for any PSR g there is a strategyproof voting rule (i.e., a mixture of strategyproof unilateral rules and duples) f_g that outputs a candidate with expected score close to the winner of g.

f) Approximate strategyproofness: We get more flexibility if on top of randomization we slightly relax the strategyproofness requirement. Two such approximations were independently suggested by Núnez and Pivato [48] and by Birrell and Pass [5]. Both solutions consider a "target rule" g, and then mix it with some carefully designed noise to obtain a randomized rule f_g that is "close" to g and "almost strategproof," where the formal meaning of these notions differ between the models. The Núnez and Pivato model makes explicit assumptions on the distribution of preferences, whereas the one by Birrell and Pass uses tools from differential privacy.

V. Voting and Mechanism Design

In contrast to the common abstract model of voting, in many specific situations agents have well-defined cardinal utilities for each alternative, and there is a clear social goal. For example - to maximize the sum of utilities. Denote the "optimal" candidate by $a^* = \text{argmax}_{a \in A} \sum_{i \in N} U_i(a)$. A natural question is then whether we can design a strategyproof mechanism that obtains or at least approximates the maximal social welfare.

A. Payments

Adding payments allows us to transfer utility between agents with much flexibility, thereby aligning their incentives.

Suppose that each voter has cardinal utilities U_i over candidates. The Vickrey-Clarke-Groves (VCG) mechanism [70],

[9], [25] collect all utility functions. Then it computes the optimal outcome a^*, 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 and a range of applications see [47].

The VCG mechanism is truthful (it is a dominant strategy for each voter to report her true utilities), and by definition it maximizes the social welfare. Both properties rely heavily on the assumption that voters' utilities are quasi-linear. Relaxing the assumption of quasi-linear utilities even slightly, leads to an impossibility result in the spirit of the Gibbard-Satterthwaite theorem [31].

B. Range Voting

Range Voting allows voters to express their cardinal preferences over candidates (normalized such that $\min_{a\in A} U_i(a) = 0$ and $\max_{a\in A} U_i(a) = 1$), and selects the one maximizing the sum of utilities. I.e. it always returns a^*.

Even without the Gibbard-Sattethwaite theorem, it is obvious that Range Voting is not truthful, as voters can always gain by reporting more extreme preferences. The G-S theorem an an even more negative implication, namely that no deterministic strategyproof mechanism can *approximate* the optimal social welfare by a factor that is sublinear in n.

Filos-Rastikas and Miltersen [20] suggested to find among the class of randomized strategyproof rules (see Sec. IV-C), the ones that give the best possible approximation for the social welfare. Their main result is a tight-to-a-constant approximation bound, that does not depend on n at all, and decreases sub-linearly with m.

C. 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 $\langle \mathcal{X}, d\rangle$ where d is a metric over the set $\mathcal{X}$. Each agent is assumed to prefer the facility to be as close as possible to her location, thus instead of reporting her entire utility function U_i, she only needs to report her location (say, some point $x_i \in \mathbb{R}^k$ or some vertex of a graph G). The cost (negative utility) of every alternative $a \in A \subseteq \mathcal{X}$ is exactly the distance $d(x_i, a)$.

The *optimal location* $a^* \in A$ is the one minimizing the *social cost* $SC(a, \boldsymbol{x}) = \sum_{i\in N} d(x_i, a)$.

A facility location mechanism is a function $g : \mathcal{X}^n \to A$, mapping the positions of all n agents to a single winning position. The special case where $A = \mathcal{X}$ is called the *unconstrained case*, as the facility can be placed anywhere, and in particular wherever an agent can be placed. Thus the contrained case is more difficult in general. The *cost approximation ratio* of g is the smallest γ s.t. for any input $\boldsymbol{x}$, $E[SC(g(\boldsymbol{x}), \boldsymbol{x})] \leq \gamma \cdot SC(a^*, \boldsymbol{x})$.

Without any restriction on the possible locations of the agents, the impossibility results of general voting rules [22], [61], [23] apply for the constrained facility location problem, which only allows for dictatorial or similar mechanisms. The welfare approximation ratio of such mechanisms can be analyzed, and shown to be $2n-1$ and $3-\frac{2}{n}$ in the deterministic and randomized cases, respectively [38].

Several papers examine variations of the problem, and in particular consider metric spaces $\mathcal{X}$ of specific shapes such as a line or a circle [62], [56], [15], [18]. For example, the Median mechanism we saw in Section IV-A provides an optimal solution for the *unconstrained* problem on a line, as agents' utilities are single-peaked.

VI. Rational Equilibrium

Once we accept that voters are going to behave strategically, and think of voting rules (with preferences) as games, we can analyze them with game theoretic tools like any other game. I will next mention several such models, which differ in their modeling assumptions.

A. Implementation

Consider any (non-resolute) SCC F, i.e. a function that maps strict preference profiles to a possibly empty set of outcomes. In what follows, $F(\boldsymbol{L}) \subseteq A$ can be thought of as some set of socially desirable alternatives under preference profile $\boldsymbol{L}$. Some examples of SCCs we might want to implement are: all Pareto optimal alternatives in $\boldsymbol{L}$; all Borda winners of $\boldsymbol{L}$ (before tie-breaking); all Condorcet winners of $\boldsymbol{L}$ (which may be empty); and so on.

Implementation of F by a mechanism g means that given any (strict) profile $\boldsymbol{L}$, a candidate $c \in A$ is in $F(\boldsymbol{L})$ if an only if voters with preferences $\boldsymbol{L}$ elect c in some *equilibrium* of g.

A most natural question is *which voting rules implement themselves* under some behavior, and if such rules even exist. This question can be extended by allowing arbitrary mechanisms that are not necessarily voting rules, and ask if a voting rule f can be implemented by some mechanism g_f using some notion of equilibrium. For example, a *truthful voting rule* implements itself in dominant strategy equilibrium.

We provide two examples of results in implementation theory that use Nash equilibrium and strong equilibrium. There are many other notions of equilibrium used in the implementation literature. Some such notions are implementation in *undominated strategies* [7], [26], *undominated Nash equilibria* [54], [64], and *mixed Nash equilibria* [26], [40].

g) NE implementation: Denote by $NE_g(\boldsymbol{L}) \subseteq A$ the set of all candidates that win in *some* Nash equilibrium of g for preferences $\boldsymbol{L}$. A mechanism g *implements a SCC F in NE*, if $NE_g(\boldsymbol{L}) = F(\boldsymbol{L})$ for all $\boldsymbol{L} \in \mathcal{L}(A)^n$.

Theorem 5 (Maskin [32])**.** *No voting rule except dictatorships and duples can be implemented in NE by any mechanism.*

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, and characterized such SCCs. A trivial example is the SCC $F(\boldsymbol{L}) = A$, which can clearly be implemented (e.g. by Plurality with $n \geq 3$ voters). A less trivial example is $F(\boldsymbol{L})$ which returns all Pareto-optimal outcomes of $\boldsymbol{L}$.

h) SE implementation: Let $SE_g(\boldsymbol{L}) \subseteq A$ 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 $G : \mathcal{L}(A)^n \rightarrow 2^A$ in SE, if $SE_g(\boldsymbol{L}) = G(\boldsymbol{L})$ for all $\boldsymbol{L}$.

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 $\boldsymbol{L}$.

Theorem 6 (Sertel and Sanver [63])**.** *The Plurality voting rule implements G_{CON} in SE, for all odd n.*

B. Bayesian uncertainty in Voting

We have seen that when voters know exactly how others are going to vote, they rarely influence the outcome. Yet it is known that people often do vote strategically, or at least trying to [57]. 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.

The classic game-theoretic approach for *games with partial information*, assumes that each player's *type* (preferences/utilities) 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 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.

Such models have been applied in a series of papers to Plurality, mainly in an attempt to explain the *Duverger Law*, which observes that in equilibrium typically only two candidates get almost all votes [60], [11], [12], [52], [19]. A general version of the model that applies for all scoring rules was suggested by Myerson and Weber [45].

An equilibrium in these models for a particular population (a distribution p_u over utility profiles) is composed of "strategies" (a mapping v from types to ballots) and "outcomes" (a distribution p_s on candidates scores), with the following self-sustaining properties: sampling a profile from p_u and apply voters' strategies v results in scores distributing according to p_s; and given p_s, voters of each type are maximizing their expected utility by voting according to v.

Other models assume voters have (uncertain) information not about other voters' preferences, but about their *actions*. Thus a voting profile is stable if every voter would choose to keep her vote assuming the outcome will be according to current profile with some "noise." This approach is highly related to election polls. Some representatives of this approach are *robust equilibrium* [39] that assumes a small probability that voters fail to cast their vote; *expectationally stable equilibrium* [53] that assumes the actual outcome an be anywhere within some small distance from the expected outcome; and *sampling equilibrium* [51] that assumes each voter is exposed to a small random sample of the other voters.

C. Other Equilibrium Models

There are many more equilibrium models, which typically focus on a specific voting rule or a class of voting rules, and make different assumptions on voters' behavior.

Some such model attribute some small positive cost to casting a ballot ("lazy bias") [13] or to casting a manipulative ballot ("truth bias") [67], [50]. This cost is sufficient to eliminate many of the unreasonable Nash equilibria of common voting rules. Another approach assumes that voters rule out dominated strategies in an iterative way until no strategy can be eliminated. While Moulin [44] showed that dominance solvable voting rules could be designed, in common rules such as Plurality, dominance leads to a unique outcome only in a small class of preferences [14].

Other models assume backward induction reasoning and analyze subgame perfect equilibria. This can apply when voters sequentially vote on parts of the decision [17], [33], [72], or when they have repeated opportunities to offer new candidates that will compete with the current winner [1].

VII. Iterative Voting and Heursitics

In the iterative voting model [37], [36], 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.

Major questions regarding iterative voting are whether it converges, and how good is the outcome to the society. For Plurality, Meir at al. [37] showed that if each voter in turn plays a best-response to the current votes of others, the game will always converge to a Nash equilibrium. This question was later studied in other voting rules, from which only Veto seems to have similar properties [59], [30], [28].

Interestingly, simulations show that in practice all common voting rules almost always converge. Further, the equilibrium result is often *better for the society* than the truthful outcome [35], [28].

A. Heursitics

Many of the objections raised in the previous sections regarding Nash equilibria, equally apply to any model based on best response, as each particular voter is unlikely to be pivotal. Things become more complicated when voters are assumed to employ more flexible *heuristics*. Such heuristics may be used either if voters only know part of the current state, or are uncertain how well it will reflect the final votes.

Some heuristics are based on common sense reasoning like focusing on the few candidates with highest current scores [58], [24], [49]. Simulations show that these heuristics almost always lead to convergence when applied in an iterative voting setting [24], [49]. Some heuristics are tailored for a specific voting rule, such are Laslier's *Leader Rule* for Approval [29], where a voter approves all candidates strictly preferred to the leader, and approves the leader if it is preferred to the runnerup.

Another approach is to generate heuristics based on dominance relations, by explicitly defining the information sets of the voters under uncertainty [10], [68], [35], [16]. One specific model is *local dominance* [35], [34], where the voter assumes that the actual candidate scores are within certain distance from the poll or current scores. In an iterative Plurality voting, this provably leads to convergence.

Extensive simulations of Plurality voting with Local Dominance heuristics show that a moderate level of uncertainty leads to the highest amount of strategic behavior, in particular when the population is diverse [35]. Further, more strategic behavior in turn leads to higher social welfare.

VIII. 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 that some models ignore or capture in different ways. Meir et al. [35] suggested a *desiderata* which is intended to provide a way to compare the strengths and weaknesses of the many different theories: some make implausible informational or cognitive assumptions, some predict behavior that is contrary to empirical evidence, some are prohibitively difficult to analyze, and some are wonderful but restricted to very specific scenarios.

I sincerely hope that familiarity with the classical and recent approaches to strategic voting will encourage researchers to develop new and better models that with improve our understanding. This, in turn, may lead to the design of improved aggregation mechanisms that lead the society to outcomes that are better for every one.

References

[1] Stéphane Airiau and Ulle Endriss. Iterated majority voting. In *Proceedings of 1st ADT*, pages 38–49, 2009.

[2] Salvador Barberà, Jordi Massó, and Alejandro Neme. Voting under constraints. *journal of economic theory*, 76(2):298–321, 1997.

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

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

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

[6] Duncan Black. On the rationale of group decision-making. *Journal of political economy*, 56(1):23–34, 1948.

[7] Tilman Börgers. Undominated strategies and coordination in normalform games. *Social Choice and Welfare*, 8(1):65–78, 1991.

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

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

[10] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *Proceedings of 25th AAAI*, 2011.

[11] Gary W Cox. Duverger's law and strategic voting. *unpublished paper, Department of Government, University of Texas, Austin*, 1987.

[12] Gary W Cox. Strategic voting equilibria under the single nontransferable vote. *American Political Science Review*, 88(3):608–621, 1994.

[13] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of 11th ACM-EC*, pages 347–356, 2010.

[14] Amrita Dhillon and Ben Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, 2004.

[15] 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.

[16] Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein. Strategic voting with incomplete information. In *Proceedings of 25th IJCAI*, 2016.

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

[18] Michal Feldman, Amos Fiat, and Iddan Golomb. On voting and facility location. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 269–286. ACM, 2016.

[19] Mark Fey. Stability and coordination in duverger's law: A formal model of preelection polls and strategic voting. *American Political Science Review*, 91(1):135–147, 1997.

[20] Aris Filos-Ratsikas and Peter Bro Miltersen. Truthful approximations to range voting. In *International Conference on Web and Internet Economics*, pages 175–188. Springer, 2014.

[21] 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.

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

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

[24] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Proceedings of 3rd ADT*, pages 181–192, 2013.

[25] Theodore Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.

[26] Matthew O Jackson. Implementation in undominated strategies: A look at bounded mechanisms. *The Review of Economic Studies*, 59(4):757–775, 1992.

[27] E. Kalai and E. Muller. Characterization of domains admitting nondictatorial social welfare functions and nonmanipulable voting procedures. *Journal of Economic Theory*, 16:457–469, 1977.

[28] Aaron Koolyk, Tyrone Strangway, Omer Lev, and Jeffrey S. Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of 26th IJCAI*, 2017. to appear.

[29] Jean-François Laslier. The leader rule: A model of strategic approval voting in a large electorate. *Journal of Theoretical Politics*, 21(1):113–136, 2009.

[30] Omer Lev and Jeffrey S Rosenschein. Convergence of iterative scoring rules. *Journal of Artificial Intelligence Research*, 57:573–591, 2016.

[31] Hongyao Ma, Reshef Meir, and David C Parkes. Social choice for agents with general utilities. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 345–351. AAAI Press, 2016.

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

[33] Richard D McKelvey and Richard G Niemi. A multistage game representation of sophisticated voting for binary procedures. *Journal of Economic Theory*, 18(1):1–22, 1978.

[34] Reshef Meir. Plurality voting under uncertainty. In *Proceedings of 29th AAAI*, pages 2103–2109, 2015.

[35] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of 15th ACM-EC*, 2014.

[36] Reshef Meir, Maria Polukarov, Jeffrey S Rosenschein, and Nicholas R Jennings. Iterative voting and acyclic games. *Artificial Intelligence*, 252:100–122, 2017.

[37] 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.

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

[39] Matthias Messner and Mattias Polborn. Robust political equilibria under plurality and runoff rule. *IGIER Working Paper*, 2005.

[40] Claudio Mezzetti and Ludovic Renou. Implementation in mixed nash equilibrium. *Journal of Economic Theory*, 147(6):2357–2375, 2012.

[41] Elchanan Mossel, Ariel D Procaccia, and Miklós Z Rácz. A smooth transition from powerlessness to absolute power. *Journal of Artificial Intelligence Research*, pages 923–951, 2013.

[42] 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.

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

[44] Hervé Moulin. Dominance solvable voting schemes. *Econometrica: Journal of the Econometric Society*, pages 1337–1351, 1979.

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

[46] Klaus Nehring and Clemens Puppe. The structure of strategy-proof social choice – part I: General characterization and possibility results on median spaces. *Journal of Economic Theory*, 135(1):269 – 305, 2007.

[47] 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.

[48] Matıas Núnez and Marcus Pivato. Truth-revealing voting rules for large populations. In *Proceedings of 6th COMSOC*. 2016.

[49] 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 29th AAAI*, pages 993–999, 2015.

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

[51] Martin J Osborne and Ariel Rubinstein. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 45(2):434–441, 2003.

[52] Thomas R Palfrey. A mathematical proof of duverger's law. 1988.

[53] Thomas R Palfrey and Howard Rosenthal. Testing game-theoretic models of free riding: New evidence on probability bias and learning. 1990.

[54] Thomas R Palfrey and Sanjay Srivastava. Nash implementation using undominated strategies. *Econometrica: Journal of the Econometric Society*, pages 479–501, 1991.

[55] Ariel D Procaccia. Can approximation circumvent gibbard-satterthwaite? In *Proceedings of 24th AAAI*, 2010.

[56] Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 177–186. ACM, 2009.

[57] Michel Regenwetter. *Behavioral social choice: probabilistic models, statistical inference, and applications*. Cambridge University Press, 2006.

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

[59] Reyhaneh Reyhani and Mark C Wilson. Best-reply dynamics for scoring rules. In *Proceedings of 20th ECAI*, 2012.

[60] William H Riker. The number of political parties: A reexamination of duverger's law. *Comparative Politics*, 9(1):93–106, 1976.

[61] 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.

[62] J. Schummer and R. V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405–428, 2004.

[63] 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.

[64] Tomas Sjöström. Implementation in undominated nash equilibria without integer games. *Games and Economic Behavior*, 6(3):502–511, 1994.

[65] Arkadii Slinko and Shaun White. Is it ever safe to vote strategically? *Social Choice and Welfare*, 43(2):403–427, 2014.

[66] 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.

[67] 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.

[68] Hans van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. Strategic voting and the logic of knowledge. In *Proceedings of 14thTARK*, 2013.

[69] Jan Van Leeuwen. *Handbook of theoretical computer science (vol. A): algorithms and complexity*. Mit Press, 1991.

[70] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[71] Lirong Xia and Vincent Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 99–108. ACM, 2008.

[72] Lirong Xia, Vincent Conitzer, and Jérôme Lang. Strategic sequential voting in multi-issue domains and multiple-election paradoxes. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 179–188. ACM, 2011.