Some Comments on Strategic Voting
Extended Abstract

Eric Pacuit* Rohit Parikh†

June 13, 2006

Abstract

Whether made explicit or implicit, knowledge theoretic properties such as common knowledge of rationality are important in understanding and modeling game-theoretic, or strategic, situations. There is a large literature devoted to exploring these and other issues related to the epistemic foundations of game theory. Much of the literature focuses on what the agents need to know about the other agents' strategies, rationality or knowledge in order to guarantee that a particular solution concept, such as the Nash equilibrium, is realized.

This paper, which is based on two recent papers[1] [7] and [16], develops a framework that looks at similar issues relevant to the field of voting theory. Our analysis suggests that an agent must possess information about the other agents' preferences in order for the agent to decide to vote strategically. In a sense, our claim is that the agents need a certain amount of information in order for the Gibbard-Satterthwaite theorem to be "effective".

1 Introduction

A comprehensive theory of multi-agent interactions must pay attention to results in social choice theory such as the Arrow and Gibbard-Satterthwaite

*Contact Author: ILLC, University of Amsterdam, epacuit@staff.science.uva.nl
†Brooklyn College and CUNY Graduate Center, rparikh@gc.cuny.edu
[1] The first paper is joint with Samir Chopra.

theorems [1, 13, 18]. These impossibility results constrain the existence of rational collective decision making procedures. In this paper we turn our attention to another aspect of social aggregation scenarios: the role played by the states of knowledge of the agents. The study of strategic interactions in game theory reflects the importance of states of knowledge of the players. In this paper, we bring these three issues—states of knowledge, strategic interaction and social aggregation operations—together.

The Gibbard-Satterthwaite theorem is best explained as follows. Let S be a social choice function whose domain is an n-tuple of preferences $P_1 \dots P_n$, where $\{1, ..., n\}$ are the voters, $\mathcal{O}$ is the set of choices or candidates and each P_i is a linear order over $\mathcal{O}$. S takes $P_1 \dots P_n$ as input and produces some element of $\mathcal{O}$ - the winner. Then the theorem says that there must be situations where it 'profits' a voter to vote *strategically*. Specifically, if P denotes the actual preference ordering of voter i, Y denotes the profile consisting of the preference orderings of all the other voters then the theorem says that there must exist P, Y, P' such that $S(P', Y) >_P S(P, Y)$. Here $>_P$ indicates: better according to P. Thus in the situation where the voter's actual ordering is P and all the orderings of the other voters (together) are Y then voter i is better off saying its ordering is P' rather than what it actually is, namely P. In particular, if the vote consists of voting for the highest element of the preference ordering, it should vote for the different highest element of P' rather than of P.

Of course, the agent might be *forced* to express a different preference. For example, if an agent, whose preferences are $B > C > A$, is only presented C, A as choices, then the agent will pick C. This 'vote' differs from the agent's true preference, but should not be understood as 'strategizing' in the true sense.

A real-life example of strategizing was noticed in the 2000 US elections when some supporters of Ralph Nader voted for their second preference, Gore,[2] in a vain attempt to prevent the election of George W. Bush. Similar examples of strategizing have occurred in other electoral systems over the years ([4] may be consulted for further details on the application of game-theoretic concepts to voting scenarios). The Gibbard-Satterthwaite theorem points out that situations like the one pointed out above *must* arise.

[2]Surveys show that had Nader not run, 46% of those who voted for him would have voted for Gore, 23% for Bush and 31% would have abstained. Hereafter, when we refer to Nader voters we shall mean those Nader voters who did or would have voted for Gore.

What interests us are the *knowledge-theoretic properties* of the situation described above. We note that unless the voter with preference P *knows* that it should vote strategically, and how, i.e., knows that the other voters' preference is Y and that it should vote according to $P' \neq P$, the theorem is not 'effective'. That is, the theorem only applies in those situations where a certain level of knowledge exists amongst voters. Voters completely or partially ignorant about other voters' preferences would have little incentive to change their actual preference at election time. In the 2000 US elections, many Nader voters changed their votes *because* opinion polls had made it clear that Nader stood no chance of winning, and that Gore could lose as a result of their votes going to Nader.

The goal of this paper is to propose a formal model in which the effect of poll information on an agent's choice of a vote can be studied. The need for such a model was suggested by Brams and Fishburn in Chapter 7 of [3]. In particular, we are interested in formally showing how voters use poll information during an election. There is a large literature which studies strategic voting in the presence of poll information. As much of the literature is geared towards a political science audience, we only discuss the papers which are related to the goals of this paper. For a discussion of formal voting theory see [4, 5]. The discussion found in Chapter 7 of [3] has much in common with this paper and so will be discussed in more detail below. For a overview of models of strategic voting in complete information environments, see [17, 14, 15]. Taking a more "computer science" approach, [9, 10, 8] provides a series of results concerning how "hard"[3] it is to take advantage of poll information. The articles [6] and [12], which compare sequential voting to simultaneous voting, both discuss issues relevant to this work. Finally, the reader is referred to [11] for a discussion of a voting procedure, called *declared-strategy voting*, which attempts to curtail the effects of strategic voting on an election.

2 A Formal Voting Model

There is a wealth of literature on formal voting theory. This section draws upon discussions in [4, 5]. The reader is urged to consult these for further details.

[3]Here "hard" is being used technically: the results are complexity theoretic.

Let $\mathcal{O} = \{o_1, \ldots, o_m\}$ be a set of candidates, $\mathcal{A} = \{1, \ldots, n\}$ be a set of agents or voters. We assume that each voter has a preference over the elements of $\mathcal{O}$, i.e., a reflexive, transitive and connected relation on $\mathcal{O}$. For simplicity we assume that each voter's preference is strict. A voter i's *strict preference relation* on $\mathcal{O}$ will be denoted by P_i. We assume that each P_i is a complete, reflexive, transitive and anti-symmetric binary relation on $\mathcal{O}$. For two candidate $o, v \in \mathcal{O}$, we will write $o >_{P_i} v$ iff $(o, v) \in P_i$ and say that i strictly prefers o to v. Henceforth, for ease of readability we will use **Pref** to denote preferences over $\mathcal{O}$. A *preference profile* is an element of $(\mathbf{Pref})^n$.

In voting scenarios such as elections, agents are not expected to announce their actual preference relation, but rather to select a vote that 'represents' their preference. Each voter chooses a vote v, an aggregation function tallies the votes of each candidate and selects a winner (or winners if electing more than one candidate). There are two components to any voting procedure. First, the type of votes that voters can cast. For example, in *plurality voting* voters can only vote for a single candidate so votes v are simply singleton subsets of $\mathcal{O}$, whereas in *approval voting* voters select a set of candidates so votes v are any subset of $\mathcal{O}$. Following [5], given a set of $\mathcal{O}$ of candidates, let $\mathcal{B}(\mathcal{O})$ be the set of feasible votes, or *ballots*. The second component of any voting procedure is the way in which the votes are tallied to produce a winner (or winners if electing more than one candidate). We assume that the voting aggregation function will select exactly one winner, so ties are always broken[4]. Note that elements of the set $\mathcal{B}(\mathcal{O})^n$ represent votes cast by the agents. An element $\vec{v} \in \mathcal{B}(\mathcal{O})^n$ is called a *vote profile*. A tallying function $\mathsf{S} : \mathcal{B}(\mathcal{O})^n \to \mathcal{O}$ maps vote profiles to candidates.

Definition 2.1 *Let $\mathcal{A}$ be a set of n agents and $\mathcal{O}$ a set of m candidates. A* **voting procedure** *is a pair $\mathcal{V} = \langle \mathcal{B}(\mathcal{O}), \mathsf{S} \rangle$, where $\mathcal{B}(\mathcal{O})$ is a set of ballots and $\mathsf{S} : \mathcal{B}(\mathcal{O})^n \to \mathcal{O}$ is a* **tallying function***, or a* **scoring function***.*

The following are examples of some well-known voting procedures. Let $\mathcal{A}$ be a set of n agents and $\mathcal{O}$ a set of m candidates.

Plurality Voting: The voting procedure $\mathcal{V}_P = \langle \mathcal{B}(\mathcal{O}), \mathsf{S} \rangle$ is called plurality voting if $\mathcal{B}(\mathcal{O}) = \{\{o\} \mid o \in \mathcal{O}\}$ and S selects the candidate with the largest number of votes. For simplicity, in the case of ties, we assume that S randomly selects among the candidates with the most votes. We assume this

[4][2] shows that the Gibbard-Satterthwaite theorem holds when ties are permitted.

throughout the paper.

Approval Voting: The voting procedure $\mathcal{V}_A = \langle \mathcal{B}(\mathcal{O}), \mathsf{S} \rangle$ is called approval voting if $\mathcal{B}(\mathcal{O}) = 2^{\mathcal{O}}$ and S selects a candidate with the largest number of approvals.

Borda Count: The voting procedure $\mathcal{V}_B = \langle \mathcal{B}(\mathcal{O}), \mathsf{S} \rangle$ is called Borda count if $\mathcal{B}(\mathcal{O}) = \mathbf{Pref}$, i.e., ballots are linear orderings of $\mathcal{O}$. The scoring function S is slightly more complicated then above. Each candidate ranked highest by a voter receives the most points, the next-highest receives the next-most points, and so on. Then S selects the candidate with the largest point total. When there are m candidates, then the usual Borda points are $m-1, m-2, \ldots, 0$ for the first choice, second choice, …, last choice.

Hare System: The voting procedure $\mathcal{V}_H = \langle \mathcal{B}(\mathcal{O}), \mathsf{S} \rangle$ is called the Hare system, or single transferable vote, if $\mathcal{B}(\mathcal{O}) = \mathbf{Pref}$ and S works as follows. If no candidate receives a majority of first-place votes, then the candidate with the fewest first-place votes is dropped and his second place votes are given to the remaining candidates. This elimination process continues until one candidate receives a simple majority.

Given a voting procedure $\mathcal{V}$ and an agent i's preference P_i, we can ask if a vote $v \in \mathcal{B}(\mathcal{O})$ is a "sincere" representation of P_i. For some voting procedures there is an objective answer to this question. For example, if we assume that the voting procedure is $\mathcal{V}_P$, then a vote v is sincere with respect to preference P iff v is the maximal[5] element of P. However, for some voting procedures, such as approval voting, more information is needed to determine whether or not a vote is a sincere reflection of a preference P. In approval voting, whether a vote v (which is a subset of $\mathcal{O}$) is sincere depends on both a preference P *and* where the agent places its cut-off point between approved and 'dis-approved' candidates.

In order to capture the above notion of a "sincere vote", we assume for each agent i a function S_i, called the **sincere vote** function, between the set of preferences $\mathbf{Pref}$ and the set of subsets of ballots. I.e., $S_i : \mathbf{Pref} \to 2^{\mathcal{B}(\mathcal{O})}$. Typically, we will assume that for each $P \in \mathbf{Pref}$, $S_i(P)$ is a singleton, but this is not necessary. If i's preference is P_i and $v \in S_i(P_i)$, then v is said

[5] Recall that we are assuming preferences are linear orders.

to be a **sincere** vote corresponding to P_i. The voter i is said to **strategize** with respect to a preference P if i selects a vote v that is not in the set $S_i(P)$. See [5] for a definition of a "sincere vote" in a variety of context and [16] for a discussion of similar issues.

Assume that the agents' true preferences are $\vec{\mathcal{P}}^* = (P_1^*, \dots, P_n^*)$ and *fixed* for the remaining discussion. Given a vote profile $\vec{v}$ of *actual* votes, we ask whether agent i will change its vote if given another chance to vote. Let $\vec{v}_{-i}$ be the vector of all *other* agents' votes. Then given $\vec{v}_{-i}$ and i's true preference P_i^*, there will be a (nonempty) set X_i of votes that are i's *best response* to $\vec{v}_{-i}$. Of course, whether v is a best response for agent i to $\vec{v}_{-i}$ will depend on the voting procedure. We will be more specific below about what constitutes a best response to a vector $\vec{v}_{-i}$ for agent i.

Suppose that for each $i \in \mathcal{A}$, f_i selects selects one such best response from X_i. We assume that agent's will only strategize if necessary. That is, if $v \in S_i(P_i^*)$ and $v \in X_i$, then f_i will select v(if more than one such v exists, then let f_i select one of these votes). Let $f(\vec{v}) = (f_1(\vec{v}_{-1}), \dots, f_n(\vec{v}_{-n}))$. We call f a **strategizing function**. If $\vec{v}$ is a fixed point of f (i.e., $f(\vec{v}) = \vec{v}$), then $\vec{v}$ is a *stable* outcome. We define f^n recursively by $f^1(\vec{v}) = f(\vec{v})$, $f^n = f(f^{n-1}(\vec{v}))$, and say that f is **stable at level** n if $f^n(\vec{v})) = f^{n-1}(\vec{v})$. It is clear that if f is stable at level n, then f is stable at all levels m where $m \geq n$. Also, if the initial votes of the $\vec{v}$ are a fixed point of f then all levels are stable.

The following two examples demonstrate the type of situations that we have in mind. The first example is taken from [3].

Example 2.2 (Brams and Fishburn [3]) *Suppose that there are four candidates $\mathcal{O} = \{o_1, o_2, o_3\}$ and nine voters divided into three groups: A, B and C. Suppose that the sizes of the groups are given as follows: $|A| = 4$, $|B| = 3$, and $|C| = 2$. We assume that all the agents in each group have the same true preference and that they all vote the same way. Suppose that the voting procedure is plurality voting ($\mathcal{V}_P$). Hence for each $i \in \mathcal{A}$, $v \in S_i(P_i)$ iff v is the maximal element of P_i. The agents' true preferences are as follows:*

$$\begin{aligned}
P_A^* &= o_1 >_{P_A^*} o_3 >_{P_A^*} o_2 \\
P_B^* &= o_2 >_{P_B^*} o_3 >_{P_B^*} o_1 \\
P_C^* &= o_3 >_{P_C^*} o_1 >_{P_C^*} o_2
\end{aligned}$$

Since we assume that in the absence of additional information, the voters will vote sincerely[6]*, candidate o_1 will win an initial election with a total of 4 votes. Now, Brams and Fishburn make the following assumption about the effect of poll information on a candidates choice of vote: "After the poll, voters will adjust their voting strategies to differentiate between the top two candidates, as indicated by the poll, if they prefer one of these candidates to the other one of these choices. Given that they are not indifferent between the top two candidates in the poll, they will vote after the poll for the one of these two they prefer" [3]. Following this protocol, only the voters in group C will change their votes. Given that they prefer o_1 to o_2, group C will give their votes to candidate o_1. Thus strengthening the lead of o_1. However, note that it is candidate o_3 who is the* Condorcet candidate*, i.e., a candidate who defeats every other candidate in a pairwise contest.*

Brams and Fishburn go on to generalize this example and show that if the agents follow the protocol described above, then under plurality voting, if the Condorcet candidate is not one of the top two candidates identified by the poll, then that Condorcet candidate will always lose. In the above example, the protocol is set up so that the second round of votes is a fixed point, i.e., the voters will not change their votes a second time. The next example describes a situation in which a fixed point does not occur until round IV. The following example was first presented in [7].

Example 2.3 *Suppose that there are four candidates $\mathcal{O} = \{o_1, o_2, o_3, o_4\}$ and five groups of voters: A, B, C, D and E. Suppose that the sizes of the groups are given as follows: $|A| = 40$, $|B| = 30$, $|C| = 15$, $|D| = 8$ and $|E| = 7$. We assume that all the agents in each group have the same true preference and that they all vote the same way. Suppose that the voting procedure is plurality voting ($\mathcal{V}_P$). Hence for each $i \in \mathcal{A}$, $v \in S_i(P_i)$ iff v is the maximal element of P_i. The agents' true preferences are as follows:*

[6]See [16] for a *proof* of the fact that voting honestly is the only protocol which dominates not voting under plurality voting.

$$\begin{aligned}
P_A^* &= o_1 >_{P_A^*} o_4 >_{P_A^*} o_2 >_{P_A^*} o_3 \\
P_B^* &= o_2 >_{P_B^*} o_1 >_{P_B^*} o_3 >_{P_B^*} o_4 \\
P_C^* &= o_3 >_{P_C^*} o_2 >_{P_C^*} o_4 >_{P_C^*} o_1 \\
P_D^* &= o_4 >_{P_D^*} o_1 >_{P_D^*} o_2 >_{P_D^*} o_3 \\
P_E^* &= o_3 >_{P_E^*} o_1 >_{P_E^*} o_2 >_{P_E^*} o_4
\end{aligned}$$

We assume that the agents all use the following protocol. If the current winner is o, then agent i will switch its vote to some candidate o' provided:

1. *i prefers o' to o, formally $o' >_{P_i} o$, and*
2. *the current total for o' plus agent i's group's votes for o' is greater than the current total for o.*

By this protocol an agent (thinking only one step ahead) will only switch its vote to a candidate which is currently not the winner. Initially, we assume that the agents all report their (unique) sincere vote. The following table describes what happens if the agents use this protocol. The candidates in bold are the winner of the current election round.

Size	*Group*	*I*	*II*	*III*	*IV*
40	*A*	$\mathbf{o_1}$	o_1	o_4	$\mathbf{o_1}$
30	*B*	o_2	$\mathbf{o_2}$	$\mathbf{o_2}$	o_2
15	*C*	o_3	$\mathbf{o_2}$	$\mathbf{o_2}$	o_2
8	*D*	o_4	o_4	o_1	o_4
7	*E*	o_3	o_3	o_1	$\mathbf{o_1}$

In round I, everyone reports their top choice and o_1 is the winner. C likes o_2 better than o_1 and its own total plus B's votes for o_2 exceed the current votes for o_1. Hence by the protocol, C will change its vote to o_2. A will not change its vote in round II since its top choice is the winner. D and E also remain fixed since they do not have an alternative like o' required by the protocol. In round III, group A changes its vote to o_4 since it is preferred to the current winner (o_2) and its own votes plus D's current votes for o_4 exceed the current votes for o_2. B and C do not change their votes. For B's top choice o_2 is the current winner and as for C, they have no o' better than o_2 which satisfies condition 2). Ironically, Group D and E change their votes to o_1 since it is *preferred to the current winner is o_2 and group A is currently voting for o_1. Finally, in round IV, group A notices that E is voting for o_1 which A prefers to o_4 and so changes its votes back to o_1. The situation stabilizes with o_1 which, as it happens, is also the Condorcet winner. I.e., it is easy to check that by following the above protocol, $f((o_1, o_2, o_2, o_4, o_1)) = (o_1, o_2, o_2, o_4, o_1)$. Thus f stabilizes at stage 4.*

We are now in a position to be more specific about what constitutes a best response for an agent. In the above examples, the agents' decisions to strategize were based on a predefined protocol. Note that in both examples, the agents behaved myopically, that is the protocol only took into account information from the current round. This restriction can be relaxed to allow the agents to make decisions based on, for example, all previous rounds. The important point from both examples is that the agents' voting strategies were explained by assuming that the agents are all following a particular protocol. In general, the agents may not necessarily all follow the same protocol as we have assumed in the above example. We will now discuss some issues relevant to formalizing this notion notion of a protocol.

In general there may be a lot of reasons why an agent may decide to change its vote. Of course an agent may change its vote *because* its preferences have changed. However, this is not the phenomenon we are trying to capture in this paper. We are interested in situations in which each agent's preference is fixed and the agent is trying to decide which vote best reflects its preference given the current situation. We assume that an agent's decision to change its vote will be based on three pieces of information. The first is the agent's actual preference. The second is information about the current vote profile, called **poll** information. The third is information about the number of agents that have the same preference.

In the above example, we assumed that during each round the agents were told the total number of votes each candidate received. In general, the form of the polling information will depend on the voting protocol that is being used. For example, knowing the total number of votes each agent received will be relevant for any voting procedure that selects the winner based solely on the total number of votes the candidates receive; however, this information will be less useful when the voting method is Borda count. In the interest of concreteness, we will restrict attention to voting methods, such as approval voting or plurality voting, that select the winner based solely on the total number of votes that the candidate receives. Thus we can model the poll information as a function $\pi : \mathcal{B}(\mathcal{O}) \rightarrow \mathbb{N}$, where $\pi(v) = n$ is interpreted as n voters have selected voter v. Let Π be the set of all such functions, i.e., the set of polls. For each poll π, let $W(\pi)$ be the candidate that would win (according to S) the election if the agents vote according[7] to π. Finally, we note that each voting profile induces a poll. That is if $\vec{v}$ is a voting profile, then we can define a poll $\pi_{\vec{v}}$ as follows, for each $v \in \mathcal{B}(\mathcal{O})$, $\pi_{\vec{v}}(v) = \sum_{v_i = v} 1$.

The second piece of information that an agent i uses to decide whether or not to change its vote is the size of the group of agents that share i's true preference. Typically, a single agent changing its votes will not affect the outcome of an election. However, as in the above example, agents will change their vote in part *because* they assume that they are part of a group which has enough weight to swing an election. In the above example, this number was constant for each agent. That is we assumed that each agent knew the exact size of the set of agents that share its actual preference. However in general, this information may not be known to an agent or the agent may only have partial information about the size of the number of agents that share its actual preference. This will be modeled by a group size function γ_i from a finite sequence of polls to the natural numbers. That is $\gamma_i : \Pi^* \rightarrow \mathbb{N}$ where $\gamma_i(\pi_1 \cdots \pi_k) = l$ means that after the series of polls $\pi_1, \ldots, \pi_n$ agent i believes that there are l agents that have the same actual preference as itself. Let Γ_i be the set of all possible such functions.

We are now in a position to formally define a protocol for an agent i. By an **election** we mean a sequence of polls, i.e., an element of Π^*. A **protocol** for agent i is a function $\Delta_i : \mathbf{Pref} \times \Pi^* \times \Gamma_i \rightarrow \mathcal{B}(\mathcal{O})$. Thus if $\sigma \in \Pi^*$ is an election and $\gamma_i \in \Gamma_i$ is a group function, then $\Delta_i(P, \sigma, \gamma_i) = v'$ means that agent i will use ballot v' in the next poll. Notice that we are assuming that the agent's use the entire election when making its decision to change its vote. This is assumed in the interest of generality. In other words, we assume that all agents have access to the election information, but whether or not they use all of that information is another issue all together. Call the vector $\vec{\Delta}$ a *group protocol*. We assume that in the absence of information, agents will vote according to their actual preferences. That is if σ is the empty string, then for each $i \in \mathcal{A}$, $\Delta_i(P_i^*, \sigma, \gamma_i) = v \in S_i(P_i^*)$.

Given a group protocol, we say that a strategizing function f is generated

[7]Of course, a poll does not list which agent voted for which candidate; however a winner can still be determined provided we assume that the S function is invariant under permutation of voters. This certainly true of many voting procedures.

by $\vec{\Delta}$, written $f_{\vec{\Delta}}$ if $f_{\vec{\Delta}}$ is defined as follows. Suppose that $\sigma = \pi_1 \cdots \pi_k$ is the current poll information and $\vec{v}$ is the current vote profile. Then we define

$$f_{\vec{\Delta}}(\vec{v}) = (\Delta_1(P_1^*, \sigma, \gamma_n(\sigma)), \ldots, \Delta_n(P_n^*, \sigma, \gamma_n(\sigma)))$$

Returning to our example. We will now demonstrate how to formalize our second example using the above machinery. First of all we assume that the agents all knew the exact number of agents that share their actual preferences. Thus for each $i \in \mathcal{A}$, γ_i is the constant function described in the example above. For example, for each election σ, $\gamma_i(\sigma) = 40$ iff $i \in A$. The Δ_i functions can be described as follows. In order to ease exposition, we will identify singleton subsets with their element. In other words, we are assuming that $\mathcal{B}(\mathcal{O}) = \mathcal{O}$. Let $W_2(\pi)$ denote the candidate that receives the second highest number of votes. The protocol that each agent follows in the first example can be described as follows. First of all we need some notation: for each pair of candidate $o, o' \in \mathcal{O}$, let $C_{P_i}(o, o')$ choose that candidate preferred according to P_i.

$$\Delta_i^1(P_i^*, \sigma, \gamma_i) = \begin{cases} o' & (v \neq W(\pi) \text{ or } v \neq W_2(\pi)) \text{ and } o' = C_{P_i^*}(W(\pi), W_2(\pi)) \\ v & \text{otherwise} \end{cases}$$

where v is the current vote (similarly for the next example). The second example can be formalized as follows:

$$\Delta_i^2(P_i^*, \sigma, \gamma_i) = \begin{cases} o' & P_i^*(o', W(\mathsf{last}(\sigma))) \text{ and } \gamma_i(\sigma) + \mathsf{last}(\sigma)(o') > \mathsf{last}(\sigma)\pi(o) \\ v & \text{otherwise} \end{cases}$$

Notice that in Δ_i^1, the group function is not used in the definition. In other words, in Example 2.2, the agents need not have any information about the size of the group that share their preferences in order to follow the protocol. Whereas, in the second example, the size of the group plays a key role in the agent's decision to change its current vote. Putting everything together, we can now define a voting model.

Definition 2.4 (Voting Model) *Given a set of n agents $\mathcal{A}$ and m candidates $\mathcal{O}$, a* **voting model** *is a tuple $\langle \mathcal{V}, \vec{\mathcal{P}}^*, \{S_i\}_{i\in\mathcal{A}}, f\rangle$ where $\mathcal{V}$ is a voting procedure, each S_i is a sincere vote function for agent i; and f is a strategizing function. We say f is generated by a group protocol $\vec{\Delta}$ if $f = f_{\vec{\Delta}}$.*

Since our candidate and agent sets are finite, if f does not stabilize then f **cycles**. We say that f has a cycle of length n if there are n different votes $\vec{\mathcal{P}}_1, \ldots \vec{\mathcal{P}}_n$ such that $f(\vec{\mathcal{P}}_i) = \vec{\mathcal{P}}_{i+1}$ for all $1 \leq i \leq n-1$ and $f(\vec{\mathcal{P}}_n) = \vec{\mathcal{P}}_1$. The following is an example of a situation in which the associated strategizing function never stabilizes:

Example 2.5 *Consider three candidates $\{o_1, o_2, o_3\}$ and 100 agents. Suppose that their are three groups of agents A, B and C. The size of each group is $|A| = 40$, $|B| = 30$ and $|C| = 30$. The actual preferences are given as follows:*

$$\begin{aligned} P_A^* &= o_1 >_{P_A^*} o_2 >_{P_A^*} o_3 \\ P_B^* &= o_2 >_{P_B^*} o_3 >_{P_B^*} o_1 \\ P_C^* &= o_3 >_{P_C^*} o_1 >_{P_C^*} o_2 \end{aligned}$$

Assume that the agents use the following protocol. An agent i will switch its vote for o to o' provided (assume w is the current winner)

1. *o' is i's second choice and the current winner is i's last choice, or*
2. *o' is i's top choice and the current winner is i's top choice.*

Assuming that the voting protocol is plurality voting and that all agents follow the above protocol generates the following table.

Size	*Group*	*I*	*II*	*III*	*IV*	*V*	*VI*	*VII*	*VIII*	*IX*	$\cdots$
40	*A*	$\mathbf{o_1}$	o_1	o_2	$\mathbf{o_2}$	$\mathbf{o_2}$	$\mathbf{o_1}$	$\mathbf{o_1}$	o_2	o_1	$\cdots$
30	*B*	o_2	$\mathbf{o_3}$	$\mathbf{o_3}$	$\mathbf{o_2}$	$\mathbf{o_2}$	o_2	o_3	$\mathbf{o_3}$	$\mathbf{o_3}$	$\cdots$
30	*C*	o_3	$\mathbf{o_3}$	$\mathbf{o_3}$	o_3	o_1	$\mathbf{o_1}$	$\mathbf{o_1}$	$\mathbf{o_3}$	$\mathbf{o_3}$	$\cdots$

After reporting their initial preferences, candidate o_1 will be the winner with 40 votes. The members of group B dislike o_1 the most, and will strategize in the next election by reporting o_3 as their preference. So, in the second round, o_3 will win. But now, members of group A will report o_2 as their preference, in an attempt to draw support away from their lowest ranked candidate. o_3 will still win the third election, but by changing their preferences (and making them public) group A sends a signal to group B that it should report its true preference - this will enable group A to have its second preferred candidate o_2 come out winner. This cycling will continue indefinitely; o_2 will win for two rounds, then o_1 for two rounds, then o_3 for two, etc.

3 Conclusion and Further Work

We have explored some properties of strategic voting and noted that the Gibbard-Satterthwaite theorem only applies in those situations where agents can obtain the appropriate knowledge. In example 5.2.2 the Condorcet winner - the winner in pairwise head-to-head contests - was picked via strategizing. Since our framework makes it possible to view opinion polls as the $n-1$ stages of an n-stage election, it implies that communication of voters' preferences and the results of opinion polls can play an important role in ensuring rational outcomes to elections. Put another way, while the Gibbard-Satterthwaite theorem implies that we are stuck with voting mechanisms susceptible to strategizing, our work indicates ways for voters to avoid irrational outcomes using such mechanisms.

References

[1] K. J. Arrow. *Social choice and individual values (2nd edition)*. Wiley, New York, 1963.

[2] Jean-Pierre Benoit. Strategic manipulation in games when lotteries and ties are permitted. *Journal of Economic Theory*, 102:421–436, 2002.

[3] Steven Brams and Peter Fishburn. *Approval Voting*. Birkhauser, Boston, 1983.

[4] Steven J. Brams. Voting Procedures. In *Handbook of Game Theory*, volume 2, pages 1055–1089. Elsevier, 1994.

[5] Steven J. Brams and Peter C. Fishburn. Voting Procedures. In *Handbook of Social Choice and Welfare*. North-Holland, 1994.

[6] S. Callander. Bandwagons and momentum in sequential voting. Working Paper: available at http://www.kellogg.northwestern.edu/faculty/callander/, 2003.

[7] Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. In Jose Julio Alferes and Joao Leite, editors, *Proceedings of JELIA 2004*, Lecture Notes in Artificial Intelligence, pages 18–30. Springer, 2004.

[8] V. Conitzer, J. Lang, and T. Sandholm. How many candidates are needed to make elections hard to manipulate? In *Proceedings of the Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2003.

[9] V. Conitzer and T. Sandholm. Vote elicitation: Complexity and strategy-proofness. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2002.

[10] V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of the International Joint Conference on Artificial Intelligences (IJCAI)*. 2003.

[11] Lorrie Faith Cranor. *Declared-Strategy Voting: An Instrument for Group Decision-Making*. PhD thesis, Washington University Sever Institute of Technonology, 1996.

[12] E. Dekel and M. Piccione. Sequential voting procedures in symmetric binary elections. *Jouranl of Political Economy*, 2000.

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

[14] Palfrey and Rosenthal. A strategic calculus of voting. *Public Choice*, 41:7 – 53, 1983.

[15] Palfrey and Rosenthal. Voter participation and strategic uncertainty. *American Political Science Review*, 41:62 – 78, 1985.

[16] Rohit Parikh and Eric Pacuit. Safe votes, sincere votes, and strategizing. Working Paper, available at staff.science.uva.nl/∼epacuit.

[17] W. Riker and P. Ordeshook. A theory of the calculus of voting. *American Political Science Review*, 62:25 – 42, 1968.

[18] Mark Satterthwaite. *The Existence of a Strategy Proof Voting Procedure: a Topic in Social Choice Theory*. PhD thesis, University of Wisconson, 1973.