



The Wisdom of Strategic Voting

QISHEN HAN, Rensselaer Polytechnic Institute, USA

GRANT SCHOENEBECK, University of Michigan, USA

BIAOSHUAI TAO, Shanghai Jiao Tong University, China

LIRONG XIA, Rensselaer Polytechnic Institute, USA

We study the voting game where agents' preferences are endogenously decided by the information they receive, and they can collaborate in a group. We show that strategic voting behaviors have a positive impact on leading to the "correct" decision, outperforming the common non-strategic behavior of informative voting and sincere voting. Our results give merit to strategic voting for making good decisions.

To this end, we investigate a natural model, where voters' preferences between two alternatives depend on a discrete state variable that is not directly observable. Each voter receives a private signal that is correlated with the state variable. We reveal a surprising equilibrium between a strategy profile being a strong equilibrium and leading to the decision favored by the majority of agents conditioned on them knowing the ground truth (referred to as the *informed majority decision*): as the size of the vote goes to infinity, every ϵ -strong Bayes Nash Equilibrium with ϵ converging to 0 formed by strategic agents leads to the informed majority decision with probability converging to 1. On the other hand, we show that informative voting leads to the informed majority decision only under unbiased instances, and sincere voting leads to the informed majority decision only when it also forms an equilibrium.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**; • **Computing methodologies** → **Distributed artificial intelligence**.

Additional Key Words and Phrases: computational social choice, information aggregation, strategic voting

ACM Reference Format:

Qishen Han, Grant Schoenebeck, Biaoshuai Tao, and Lirong Xia. 2023. The Wisdom of Strategic Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*, July 9–12, 2023, London, United Kingdom. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3580507.3597681>

1 INTRODUCTION

Today, voting is used to make an array of binary decisions permeating nearly every corner of life including in recall/run-off public elections, adoption of decrees by religious institutions, decisions by corporate boards on whether or not to pursue a new strategy/acquisition/etc, hiring and by-law decisions at university, and public entertainments like talent shows. In most cases, the voting is attempting to aggregate both the agents' preferences and knowledge. A key aspect of this setting is that agents have preferences over outcomes contingent on some underlying state that they cannot directly observe, and the goal is to make a "good" decision that reflects the preferences of the agents.

A full version of this paper is available at <https://arxiv.org/abs/2305.11021>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

EC '23, July 9–12, 2023, London, United Kingdom

© 2023 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0104-7/23/07...\$15.00

<https://doi.org/10.1145/3580507.3597681>

EXAMPLE 1. Suppose the voters vote to decide the policy towards the COVID-19 pandemic. The two choices are to accept the more-restrictive policy (Accept) and to keep the status quo (Reject). The consequence of the policy depends on the fact that the COVID virus is of high or low risk, and more people tend to accept the policy when COVID is of high risk than when COVID is of low risk. The voters do not know the risk level of the virus directly. Instead, every voter forms a private judgment on the risk level based on his/her own information sources. voters may have different opinions on whether to accept the policy which may or may not depend on the risk level. Can the voters achieve a good decision via the majority vote?

Three different lines of work aim to address this problem under different models and with different goals. The first line of work is axiomatic social choice [Arrow 2012; Plott 1976], where agents’ preferences are exogenously given, and the goal is to design voting rules that satisfy desiderata, often called *axioms*, especially when agents *sincerely* report their preferences. The second line of work is along the extensions of the Condorcet Jury Theorem [Condorcet 1785], where agents’ preferences are endogenous and depend on the information structure and the signals they receive. The goal is to design mechanisms to reveal the true state of the world, especially when agents vote *informatively*, i.e., their votes honestly reflect the private signals they receive. The survey by Nitzan and Paroush [2017] provides a comprehensive overview. The third line of work originated from Feddersen and Pesendorfer [1997], where agents’ preferences are endogenous as in the second line of work, yet the goal is different. Instead of revealing the true state of the world, the goal is to achieve *informed majority decision*, which is the decision favored by the majority of the agents if the world state were known to them. Our paper is along the third line of work.

1.1 Strategic Behaviors

Previous work shows that a good decision can be reached when agents follow sincere or informative behaviors. However, when agents are *strategic*, they may have incentives to deviate from sincere or informative voting to achieve a preferred result with a higher probability. This is not a problem for axiomatic social choice, as strategic agents will always vote for their preferred alternative in binary voting [Barberà et al. 1991]. However, when agents have preferences decided by uncertain world states, the surprising result by Austen-Smith and Banks [1996] shows that even in binary voting, informative voting may fail to form a Nash equilibrium. The key insight is that an agent’s vote makes a difference only when all other votes form a tie, which means that when an agent strategically thinks about his/her vote, effectively he/she gains more information about the ground truth (by assuming that other votes are tied). This is illustrated in the following example.

EXAMPLE 2. Consider an instance of the COVID policy problem, where the utility of the agents and signal distribution of different risk levels are shown in the tables below.

State	High Signal	Low Signal
High Risk	0.9	0.1
Low Risk	0.4	0.6

Table 1. Signal distributions.

State	Accept	Reject
High Risk	1	0
Low Risk	0	1

Table 2. Agents’ utilities.

Suppose all but one agents are informative and the remaining agent is strategic. Informative agents vote for Accept when they receive a high signal and Reject when they receive a low signal. The strategic agent only cares about the pivotal case where exactly half of the informative agents vote for Accept. However, given that agents receive high signals with a probability of 0.9 given the risk level being high, the pivotal case implies a high probability that the risk level is low, and the strategic agent will vote for Reject even after receiving a high signal.

The above deviation of strategic binary voting with preferences endogenously affected by the unobservable world state sharply contrasts with the axiomatic social choice where preferences are exogenously given. In the latter case, the majority rule is strategy-proof while the former case attracts a large literature to study the binary voting problem under game theoretical contexts, studying the impact of strategic behavior. For the truth-revealing goal, Wit [1998] and Myerson [1998] show that a selected equilibrium with mixed strategy reveals the world state with high probability. Feddersen and Pesendorfer [1998] show the existence of such equilibrium in any non-unanimous voting, while in unanimous voting strategic voting has a constant probability to make a mistake. And for the informed majority decision, Feddersen and Pesendorfer [1997] adopt a model with continuous world states and an asymptotically large number of agents whose preferences are drawn from a distribution with full support on a continuum and show that the equilibrium is unique and always leads to the informed majority decision with high probability. Schoenebeck and Tao [2021] proposes a mechanism incentivizing informative voting from agents and leading to the informed majority decision with high probability.

Nevertheless, there are two aspects not addressed by previous works. Firstly, previous works (except for Schoenebeck and Tao [2021]) focus on Nash equilibrium which allows only individual manipulation. In real-world scenarios, on the other hand, such strategic manipulation often occurs in a *coalition* of agents. Coalitional manipulation is more powerful than individual manipulation as it allows multiple agents to coordinate and deviate at the same time. The following example shows that a Nash equilibrium is still prone to a group of manipulators in binary voting.

EXAMPLE 3. Consider an instance with three agents, whose utility is shown as follows.

Agent	(High, Accept)	(High, Reject)	(Low, Accept)	(Low, Reject)
1	1	0	1	0
2	1	0	0	1
3	0	1	0	1

Table 3. The utility of three agents under different states and decisions.

The following strategy profile is a Nash equilibrium: agent 1 always votes for Accept, agent 2 votes informatively, and agent 3 always votes for Reject. Here, agent 1 and 3 play their dominant strategies, and, consequently, informative voting is the best strategy for agent 2. However, this strategy profile is dominated by the profile where all agents vote informatively. Under informative voting, the decision in accord with the state (Accept in High state, and Reject in Low state) is selected with a larger probability, and the utility of agent 2 increases. On the other hand, the overall probability of choosing to Accept or Reject does not change, so agent 1 and 3’s utilities remain the same.

Secondly, previous works focus on the existence of certain equilibria that achieves the goal (revealing the world state or reaching the informed majority decision). However, the existence of multiple equilibria [Wit 1998], including “bad equilibria” that do not lead to the goal, makes the behavior of strategic agents unpredictable, as it is uncertain which equilibrium agents will play. One response to multiple equilibria is to select an equilibrium that is more “natural” or “reasonable” than others, named *equilibrium selection*. However, equilibrium selection cannot guarantee that agents will play the selected equilibrium, as it is unclear which equilibrium is more “natural” or “reasonable” in many scenarios, and agents may not agree on a “more natural” equilibrium even if it exists.

As a consequence, the following research question remains unanswered: **does binary voting always lead to the informed majority decision with coalitional strategic agents?**

1.2 Our contribution

We give a surprising confirmative answer to this question under mild conditions. We show that coalitional strategic behaviors positively impact achieving the informed majority decision and outperform non-strategic voting. We show that every equilibrium leads to the informed majority decision, and every voting profile that leads to the informed majority decision is an equilibrium. On the contrary, non-strategic behaviors lead to the informed majority decision only under certain conditions. Our results give merit to strategic behaviors and extend Feddersen and Pesendorfer's results to settings with coalitional strategic agents.

We study the solution concept of ε -strong Bayes Nash Equilibrium, which precludes groups of agents from reaching higher expected utilities by coordinating. We show the equivalence of a strategy profile being “good” (leading to the informed majority decision with high probability, or, equivalently, of high *fidelity*) and being an ε -strong Bayes Nash Equilibrium with $\varepsilon = o(1)$ (Theorem 2). We also guarantee the existence of an ε -strong Bayes Nash with $\varepsilon = o(1)$ in any instance (Theorem 3).

On the other hand, we characterize the conditions where strategy profiles succeed and fail to achieve the informed majority decision. Applying these results, we study two common non-strategic behavior – *informative voting*, where agents honestly reflect their private information in their votes, and *sincere voting*, where agents vote as if they are the only decision-maker. We show that (1) informative voting leads to the informed majority decision only when the majority vote threshold is unbiased compared with the signal distribution (Corollary 3), and (2) sincere voting leads to the informed majority decision only when it is also an equilibrium (Corollary 4). These observations indicate that strategic behavior “prevails” over non-strategic behaviors in binary voting!

The technical key for the probability analysis is to compute the *excess expected vote share*, i.e., the amount of expected vote share an alternative attracts that exceeds the threshold, and to upper (or lower) bound the fidelity given different cases of excess expected vote shares. A strategy profile has high fidelity if and only if its excess expected vote share is strictly positive (Theorem 4).

We follow the setting in Schoenebeck and Tao [2021], which is an extension of the setting in Austen-Smith and Banks [1996], and consider agents with preferences contingent on underlying world states in a single framework. Also, as in Example 1 and previous work, we assume that various constraints in the real world prevent discussion after agents see their signals. Such constraints can be of a time aspect (a quick decision must be made and there is no time for discussion), a procedural aspect (a formal conference that prohibits participants from discussing privately before voting), and/or a societal aspect (it is socially unsuitable to discuss some preferences), etc. Therefore, we consider an *ex-ante* setting where the expected utilities are computed before agents receive their signals.

1.3 Related Work

The famous Jury Theorem from Condorcet [1785] has “formed the basis for the development of social choice and collective decision-making as modern research fields” [Nitzan and Paroush 2017]. The theorem states that a group of decision-makers could reveal the correct world state with a higher probability than any individual in the group, and such probability converges to 1 as the number of group members increases. A large literature on collective decision-making has followed Condorcet's path trying to extend the result into more general models [Boland et al. 1989; Grofman et al. 1983; Miller 1986; Owen et al. 1989].

The game-theoretical study of the Condorcet Jury Theorem starts from Austen-Smith and Banks [1996]. Austen-Smith and Banks study a collaborative voting game where each agent shares the same preference and receives a binary signal correlated with an unknown binary state of the

world. However, even in this case, they showed that sincere voting and informative voting do not always form a Nash Equilibrium. As a consequence, the following works focus on the effect of strategic behavior in the majority vote and propose equilibria that reveal the ground truth [Duggan and Martinelli 2001; Feddersen and Pesendorfer 1998; Meirowitz 2002; Myerson 1998; Wit 1998]. Feddersen and Pesendorfer [1997] adopt a similar information structure with the game theoretical study of Condorcet Jury Theorem but aim to achieve a different goal of informed majority decision. Other generalizations of the Condorcet Jury Theorem include dependent agents [Kaniovski 2010; Nitzan and Paroush 1984; Shapley and Grofman 1984], agents with different competencies [Ben-Yashar and Zahavi 2011; Gradstein and Nitzan 1987; Nitzan and Paroush 1980], and voting with more than two alternatives [Goertz and Maniquet 2014; Young 1988].

Another line of work related to collective decision-making focuses on designing mechanisms that lead to the correct decision. Recent work shows the reliability of the “surprisingly popular” answer when agents are sincere [Hosseini et al. 2021; Prelec et al. 2017] and strategic [Schoenebeck and Tao 2021]. In particular, Schoenebeck and Tao [2021] adopt the “surprisingly popular” technique into a social choice context with strategic agents, and propose a truthful mechanism to aggregate information. They show that even in a setting where agents have subjective preferences contingent on an objective underlying state, their mechanism reveals the informed majority decision with high probability and is an (ex-ante) ϵ -strong Bayes Nash Equilibrium with ϵ converging to 0 at an exponential rate. Our work follows the setting in Schoenebeck and Tao’s work, but our work is different in that the aggregation happens implicitly because agents are acting strategically rather than because a mechanism explicitly selects a surprisingly popular answer.

Our work is also related to *information elicitation*, which aims to collect truthful and high-quality information from agents under a noisy information structure. Information elicitation is well developed with multiple lines of research focusing on different aspects of the problem, including scoring rules [Bickel 2007; Gneiting and Raftery 2007], peer prediction mechanisms [Miller et al. 2005; Schoenebeck and Yu 2021; Schoenebeck et al. 2021], Bayesian Truth Serum [Prelec 2004; Witkowski and Parkes 2012], and prediction markets [Miller and Drexler 1988; Wolfers and Zitzewitz 2004]. Unfortunately, information elicitation is incompatible with the voting scenario in our paper for two reasons. Firstly, information elicitation requires agents to be indifferent to the outcome, while agents are incentivized by the outcome of the vote. Secondly, information elicitation uses payments to reward the agents, while voting does not have monetary rewards.

2 MODELS AND PRELIMINARIES

We first present our model and results with binary world states and binary private signals, which convey the main ideas of this work while also hiding much of the complexity. The general extension into the non-binary setting is in Section 5. We follow the setting in Schoenebeck and Tao [2021] and consider agents with subjective preferences contingent on an objective underlying state in one framework.

Alternatives and World States. N agents vote for two alternatives **A** (standing for “accept”) and **R** (standing for “reject”). There are $K = 2$ possible world states $\mathcal{W} = \{L, H\}$ (standing for “low risk” and “high risk” respectively), where **A** is more preferred in H , and **R** is more preferred in L . We use k to denote a generic world state. The world state is not directly observable by the agents. Let $P_H = \Pr[W = H]$ and $P_L = \Pr[W = L]$ be the common prior of the world states. We assume $P_H > 0$ and $P_L > 0$.

Private Signals. Every agent receives a signal in $\mathcal{S} = \{l, h\}$. We use m to denote a generic signal, and S_n to denote the random variable representing the signal that agent n receives. We assume the signals agents receive are independent and have identical distributions conditioned on the world

state. Let $P_{mk} = \Pr[S_n = m \mid W = k]$ be the probability that an agent receives signal m under world state k . The signal distributions $((P_{hH}, P_{lH}), (P_{hL}, P_{lL}))$ are also common knowledge. We assume that the signals are positively correlated to the world states. Specifically, we have $P_{hH} > P_{hL}$ and $P_{lH} < P_{lL}$. On the other hand, we allow biased signals and DO NOT assume $P_{hH} > P_{lH}$ or $P_{hL} < P_{lL}$.

Majority Vote. This paper considers the majority vote with threshold μ . Each agent n votes for A or R. If at least $\mu \cdot N$ agents vote for A, A is announced to be the winner; otherwise, R is announced to be the winner.

Utility and Types of Agents. Each agent n has a utility which is a function of the true world state and the outcome of the vote. Formally, we have $v_n : \mathcal{W} \times \{A, R\} \rightarrow \{0, 1, \dots, B\}$, where B is the positive integer upper bound. We assume that A is more preferable in H than in L , and R is the opposite: for every agent n , $v_n(H, A) > v_n(L, A)$ and $v_n(H, R) < v_n(L, R)$.

The different endogenous preferences of agents are reflected by different utility functions. *Predetermined* agents always prefer the same alternative, and *contingent* agents have preferences depending on the world state. Predetermined agents can be further divided into *friendly* and *unfriendly* agents based on the alternative they prefer. For an agent n , if n is a friendly agent, $v_n(H, A) > v_n(L, A) > v_n(L, R) > v_n(H, R)$; if n is an unfriendly agent, $v_n(L, R) > v_n(H, R) > v_n(H, A) > v_n(L, A)$; and if n is a contingent agent, $v_n(H, A) > v_n(H, R)$ and $v_n(L, R) > v_n(L, A)$.

Let α_F , α_U , and α_C be the approximated fraction of each type of agent. Formally, given N agents, $N_F = \lfloor \alpha_F \cdot N \rfloor$ is the number of friendly agents, $N_U = \lfloor \alpha_U \cdot N \rfloor$ is the number of unfriendly agents, and $N_C = N - N_F - N_U$ is the number of contingent agents. α_F , α_U , and α_C are common knowledge and do not depend on N .

Informed Majority Decision. The goal of the voting is to output the informed majority decision, which is the alternative favored by the majority of the agents if the world state were known. The informed majority decision shares the same threshold μ as the majority vote threshold. If A is preferred by at least $\mu \cdot N$ agents, then A is the informed majority decision; otherwise, R is the informed majority decision.

In this paper, we assume that neither friendly agents nor unfriendly agents can dominate the vote. Otherwise, the informed majority decision does not depend on the state and one coalition can always enact it via a dominant strategy. As a result, A is the informed majority decision when the world state is H , and R is the informed majority decision when the world state is L .

EXAMPLE 4. Consider the COVID policy-making scenario. $N = 20$ voters decide whether to accept (denoted as A) or reject (denote as R) the more-restrictive policy. The world state $\{L, H\}$ describes the real risk level of the virus. $W = H$ means high risk level, and $W = L$ means low risk level. The voters' beliefs form a common prior based on some preliminary reports. Suppose $P_H = 0.4$ and $P_L = 0.6$, which means the risk level has a prior probability of 0.4 to be high.

Every voter receives a private signal l or h from his/her information sources. The signals somehow reflect the risk level but are noisy. Suppose this is a biased scenario (for example, there has been a boost of positive cases in the past week), and members are always more likely to receive the high signal. For example, $P_{hH} = 0.8$ and $P_{hL} = 0.6$, i.e., a voter will receive an h signal with probability 0.8 if the risk level is high and receive an h signal with probability 0.6 if the risk level is low.

The majority vote threshold is $\mu = 0.6$. Therefore, A is the winner if and only if at least 12 voters vote for it. There are 4 friendly voters, 6 unfriendly voters, and 10 contingent voters. The informed majority decision depends on the world state: "accept" is the informed majority decision if the world state is H , and "reject" is the informed majority decision if the world state is L .

We assume that agents of the same type share the same utility function (which may not be true in general) shown in Table 4.

(Winner, World State)	(A, H)	(A, L)	(R, H)	(R, L)
Friendly agent	8	6	2	4
Unfriendly agent	3	1	5	8
Contingent agent	3	2	1	8

Table 4. Utility of agents in Example 4.

Strategy. A (mixed) strategy is a mapping from the agent’s private signal to a distribution on $\{\mathbf{A}, \mathbf{R}\}$. For a set S , let $\Delta(S)$ be the set of all possible distributions on S . Formally, an agent n ’s strategy $\sigma_n : \mathcal{S} \rightarrow \Delta(\{\mathbf{A}, \mathbf{R}\})$. A strategy can be represented as a vector $\sigma = (\beta_l, \beta_h)$, where β_m is the probability that the agent votes for $\mathbf{A}$ when receiving signal m . A strategy profile is the vector of strategies of all agents. $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$. We call a strategy profile Σ a *symmetric strategy profile induced by strategy σ* if all agents play the same strategy σ in Σ .

Definition 1. An informative strategy is $\sigma = (0, 1)$, i.e. voting for $\mathbf{A}$ when receiving h and voting for $\mathbf{R}$ when receiving l . A strategy profile is informative when every agent votes informatively.

In this paper, we focus on *regular* strategy profiles.

Definition 2. A strategy profile Σ is *regular* if all friendly agents always vote for $\mathbf{A}$, and all unfriendly agents always vote for $\mathbf{R}$ in Σ .

We believe this restriction is mild and natural since “always vote for $\mathbf{A}$ ” is the dominant strategy for a friendly agent, and “always vote for $\mathbf{R}$ ” is the dominant strategy for an unfriendly agent in the majority vote.

Fidelity and Expected Utility. Given a strategy profile Σ , let $\lambda_k^{\mathbf{A}}(\Sigma)$ ($\lambda_k^{\mathbf{R}}(\Sigma)$, respectively) be the (ex-ante, before agents receiving their signals) probability that $\mathbf{A}$ ($\mathbf{R}$, respectively) becomes the winner when the world state is k .

Definition 3 (Fidelity). Fidelity is the likelihood that the informed majority decision is reached. In our setting, the fidelity when agents play strategy profile Σ is

$$A(\Sigma) = P_L \cdot \lambda_L^{\mathbf{R}}(\Sigma) + P_H \cdot \lambda_H^{\mathbf{A}}(\Sigma).$$

We use the word *fidelity* to distinguish the notion from *accuracy*, which usually denotes the likelihood that the correct world state is revealed.

The (ex-ante) expected utility of an agent n exclusively depends on $\lambda_k^{\mathbf{A}}(\Sigma)$ and $\lambda_k^{\mathbf{R}}(\Sigma)$:

$$u_n(\Sigma) = P_L(\lambda_L^{\mathbf{A}}(\Sigma) \cdot v_n(L, \mathbf{A}) + \lambda_L^{\mathbf{R}}(\Sigma) \cdot v_n(L, \mathbf{R})) + P_H(\lambda_H^{\mathbf{A}}(\Sigma) \cdot v_n(H, \mathbf{A}) + \lambda_H^{\mathbf{R}}(\Sigma) \cdot v_n(H, \mathbf{R})).$$

Instance and Sequence of Strategy Profiles. We define an instance $\mathcal{I}$ of a voting game on the agent number N , the majority vote threshold μ , the world state prior distribution (P_L, P_H) , the signal distributions $((P_{hH}, P_{lH}), (P_{hL}, P_{lL}))$, the utility functions of all the agents $\{v_n\}_{n=1}^N$, and the approximated fraction of each type $(\alpha_F, \alpha_U, \alpha_C)$. Let $\{\mathcal{I}_N\}_{N=1}^\infty$ (or $\{\mathcal{I}_N\}$ for short) be a sequence of instances, where each $\mathcal{I}_N$ is an instance of N agents. The instances in a sequence share the same parameters $\{\mu, (P_L, P_H), ((P_{hH}, P_{lH}), (P_{hL}, P_{lL})), (\alpha_F, \alpha_U, \alpha_C)\}$. We do not regard agents in different instances as related and have no additional assumption on the utility functions of agents.

We define a sequence of strategy profiles $\{\Sigma_N\}_{N=1}^\infty$ on an instance sequence $\{\mathcal{I}_N\}$. Similarly, we do not have additional assumptions about the agents. Therefore, for different instances in the sequence, the strategies and utility functions of agents can be drastically different. A strategy profile sequence $\{\Sigma_N\}$ is symmetric and induced by strategy σ if every strategy profile Σ_N in the sequence

is a symmetric strategy profile induced by σ . A sequence of strategy profiles is regular if every strategy profile in the sequence is a regular profile.

ε -strong Bayes Nash Equilibrium. In this paper, we use the solution concept of ε -strong Bayes Nash Equilibrium, an approximation of strong Bayes Nash Equilibrium where no group of agents can increase their utilities by more than ε through deviation. A strategy profile $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$ is an *ε -strong Bayes Nash Equilibrium* (ε -strong BNE) if there does not exist a subset of agents D and a strategy profile $\Sigma' = (\sigma'_1, \sigma'_2, \dots, \sigma'_N)$ such that

- (1) $\sigma_n = \sigma'_n$ for all $n \notin D$;
- (2) $u_n(\Sigma') \geq u_n(\Sigma)$ for all $n \in D$; and
- (3) there exists $n \in D$ such that $u_n(\Sigma') > u_n(\Sigma) + \varepsilon$.

By definition, when $\varepsilon = 0$, the equilibrium is a strong Bayes Nash Equilibrium where no group of agents can strictly increase their utilities through deviation. Unfortunately, a strong BNE does not always exist, as shown in the following theorem. Therefore, we seek ε -strong BNE as an approximation.

THEOREM 1. *For any $N_0 \in \mathbb{N}$, there exists an instance of $N > N_0$ agents, in which a strong Bayes Nash Equilibrium does not exist.*

PROOF SKETCH. For any $N_0 \in \mathbb{N}$, we construct an instance of $N = 2N_0 + 3$ agents. The agents consist of three parts: F is a set of $N_0 + 1$ friendly agents. C is a set of two contingent agents. And U is a set of N_0 unfriendly agents. Agents in the same set share the same utility, which is shown in Table 5. The threshold is $\mu = 0.5$. The prior distribution is $P_L = P_H = 0.5$. The signal distribution is $P_{hH} = P_{lL} = 0.8$ and $P_{lH} = P_{hL} = 0.2$.

Agents	$v(H, A)$	$v(L, A)$	$v(L, R)$	$v(H, R)$
F	100	99	1	0
C	90	0	100	0
U	1	0	100	99

Table 5. Utility of three groups

Agents	Σ_1	Σ_2	Σ_3
F	50.396	66.14	50.3
C	85.12	75.2	76
U	50.396	34.46	50.3

Table 6. Expected utility under three profiles

Consider the following three strategy profiles, under which the expected utility of each group is shown in Table 6.

- Σ_1 : N_0 agents in F always vote for A , and one agent votes informatively. C vote informatively. U always vote for R .
- Σ_2 : F always vote for A . C vote informatively. U always vote for R .
- Σ_3 : F always vote for A . One C agent votes informatively, and the other always votes for R . U always vote for R .

These three strategy profiles form a cycle $\Sigma_1 \rightarrow \Sigma_2 \rightarrow \Sigma_3 \rightarrow \Sigma_1$ of deviation, where a group of agents has incentives to deviate to the next profile.

For any other strategy profile Σ , there exists a group of agents with incentives to deviate to one of the three profiles. Firstly, F agents and U agents would like to deviate from their dominant strategy of always voting for A (R , respectively) whenever it can increase the probability that their preferred candidate wins. Given F and U agents play dominant strategies, the best strategy for two C agents is to play the strategy in Σ_3 (one agent votes informatively, the other votes for R). Then we know that an F agent and two C agents have incentives to deviate from Σ_3 to Σ_1 . Therefore, there does not exist a strong Bayes Nash in this instance. Full proof is available in the full version. $\square$

3 EQUIVALENCE BETWEEN HIGH FIDELITY AND STRONG EQUILIBRIUM

In this section, we show that strategic behaviors indeed have a positive impact on leading to the informed majority decision. Theorem 2 states that if the fidelity of a regular (Definition 2) strategy profile sequence $\{\Sigma_N\}_{N=1}^\infty$, i.e., $A(\Sigma_N)$, converges to 1 as N goes to infinity, every Σ_N in the sequence will be an ε -strong Bayes Nash Equilibrium where ε converges to 0. On the other hand, if $A(\Sigma_N)$ does not converge to 1, then we can find infinitely many Σ_N that are not ε -strong BNE with a constant ε . Moreover, Theorem 3 guarantees that there always exists a regular strategy profile whose fidelity converges to 1, which leads to an ε -strong BNE with $\varepsilon = o(1)$. The two theorems together indicate that strategic voting leads to the informed majority decision in any sequence of instances.

THEOREM 2. *Given an arbitrary sequence of instances and an arbitrary regular strategy profile sequence $\{\Sigma_N\}_{N=1}^\infty$, let $\{A(\Sigma_N)\}_{N=1}^\infty$ be the sequence of the fidelities of Σ_N .*

- *If $\lim_{N \rightarrow \infty} A(\Sigma_N) = 1$, then for every N , Σ_N is an ε -strong BNE with $\varepsilon = o(1)$.*
- *If $\lim_{N \rightarrow \infty} A(\Sigma_N) = 1$ does not hold, then there exist infinitely many N such that Σ_N is NOT an ε -strong BNE for some constant ε .*

THEOREM 3. *Given any arbitrary sequence of instances, there always exists a sequence of regular strategy profiles $\{\Sigma'_N\}_{N=1}^\infty$ such that $A(\Sigma'_N)$ converges to 1.*

We first give a concrete example to illustrate Theorem 2, in which we show an instance for each case in the theorem.

EXAMPLE 5. *We follow the setting of Example 4 except for two differences. First, there is a series of $N = 20, 30, \dots, 500$. For each N , the ratio of friendly, unfriendly, and contingent agents is fixed at $2 : 3 : 5$. Second, we consider two different cases of signal distributions that fall into different cases of Theorem 2. They share the same signal distribution in world state L : $P_{lL} = 0.8, P_{hL} = 0.2$, but the signal distribution for H is different. In case (1), $P_{lH} = 0.1, P_{hH} = 0.9$; and in case (2), $P_{lH} = 0.25, P_{hH} = 0.75$.*

We focus on regular strategy profiles where all contingent agents vote informatively (Definition 1). For case (2), we also consider another series of regular strategy profiles Σ'_N where contingent agents play $\sigma' = (0.48, 0.96)$. In Example 7 and Theorem 4 later, we verify that the fidelity of the regular informative voting converges to 1 in case (1) but does not converge to 1 in case (2). On the other hand, the fidelity of the deviating strategy profile Σ'_N in case (2) converges to 1. Figure 1(a) illustrates these trends of fidelity.

The expected utilities of contingent agents in different cases and strategies are shown in Figure 1(b). Note the maximum expected utility that a contingent agent can get is $0.4 \times 3 + 0.6 \times 8 = 6.0$. In accordance with the fidelity, the expected utility of Σ_N in case (1) converges to the maximum. In case (2), on the other hand, Σ_N is dominated by Σ'_N by a utility gain of at least 0.4. Therefore, the group of contingent agents has no incentive to deviate in case (1) but has incentives to deviate to Σ'_N in case (2).

3.1 Proof Sketch of Theorem 2

To show the relationship between $A(\Sigma_N)$ and ε , we have the following lemma.

LEMMA 1. *For every N , a regular strategy profile Σ_N is an ε -strong BNE with $\varepsilon = 2B(B+1)(1 - A(\Sigma_N))$, where B is the upper bound of utility function v_n .*

Lemma 1 is an extension of Theorem 3.3 in Schoenebeck and Tao [2021]. It shows that every Σ_N is an ε -strong BNE with ε proportional to $1 - A(\Sigma_N)$. To prove Lemma 1, we show that, for any other strategy profile Σ'_N , a group of agents with incentives to deviate does not exist.




Fig. 1. Fidelity and expected utilities of informative voting.

There are two cases of Σ'_N . In the first case, the fidelity Σ'_N is bounded by the fidelity of Σ_N . More precisely, $(1 - A(\Sigma'_N)) < (B + 1) \cdot (1 - A(\Sigma_N))$. Then two profiles do not make a big difference, and no agent can gain more than $\varepsilon = 2B(B + 1)(1 - A(\Sigma_N))$ after deviation (Claim 1).

In the second case, the fidelity Σ'_N is unbounded and much worse than the fidelity of Σ_N . Then any contingent agent has no incentives to deviate, as their expected utilities are most positively correlated with fidelity (Claim 2). Next, we show that a deviating group cannot contain both friendly and unfriendly agents, because if the expected utility of one side increases by more than ε , the other side's will decrease (Claim 3). Therefore, a deviating group contains either only friendly agents or only unfriendly agents. Finally, we show that in neither case can the deviation succeed, because pre-determined agents have already played their dominant strategies in a regular profile (Claim 4).

Now we are ready to propose the proof for Theorem 2. We will actually use Lemma 1 and Theorem 3 to prove Theorem 2. We will discuss the two cases separately.

When $\lim_{N \rightarrow \infty} A(\Sigma_N) = 1$, we apply Lemma 1 to each Σ_N , and get that every Σ_N is an ε -strong BNE where $\varepsilon = 2B(B + 1) \cdot (1 - A(\Sigma_N))$. Then ε will converge to 0 as $N \rightarrow \infty$.

When $\lim_{N \rightarrow \infty} A(\Sigma_N) = 1$ does not hold, there are infinitely many N with Σ_N being of low fidelity. By Theorem 3 there exists a regular strategy profile sequence $\{\Sigma'_N\}$ with fidelity converging to 1. Because of the difference in fidelity, there are infinitely many N such that $\Sigma_N \neq \Sigma'_N$. Then we show that, for all sufficiently large N where Σ_N is of low fidelity, if all contingent agents turn to play Σ'_N from Σ_N , every contingent agent will gain at least a constant amount of extra utility. Therefore, for infinitely many N , Σ_N is NOT an ε -strong BNE for some constant ε . Full proof is available in the full version.

3.2 Proof Sketch of Theorem 3

In the proof of Theorem 3, we construct a strategy σ' and show that the regular strategy profile sequence $\{\Sigma'_N\}$ where all contingent agents play σ' has fidelity that converges to 1. It suffices to construct σ' such that

- (1) if H is the actual world state, the expected fraction of the voters voting for A is more than μ by a constant;
- (2) if L is the actual world state, the expected fraction of the voters voting for A is less than μ by a constant.

If this is true, $A(\Sigma'_N)$ converges to 1 due to the Hoeffding Inequality. It remains to construct σ' such that (1) and (2) hold.

We first construct σ'_μ such that the expected fraction of the voters voting for **A** is exactly μ , where in σ'_μ the contingent voter votes for **A** with a probability that is independent to the signal she receives. This can be done by setting $\sigma'_\mu = (\beta^*, \beta^*)$ where β^* satisfies $\alpha_F + \alpha_C \cdot \beta^* = \mu$ (notice that, given the fraction α_F of the friendly voters who always vote for **A**, the fraction α_U of the unfriendly voters who never vote for **A**, and the fraction α_C of the contingent voters who vote for **A** with probability β^* , the expected fraction of votes for **A** is $\alpha_F + \alpha_C \cdot \beta^*$).

Next, we will adjust σ'_μ to $\sigma' = (\beta_l, \beta_h)$ that satisfies (1) and (2). Naturally, we would like to increase the probability for voting **A** if an h signal is received, and we would like to decrease this probability if l is received. That is, we have $\beta_l = \beta^* - \delta_l$ and $\beta_h = \beta^* + \delta_h$ for some $\delta_l, \delta_h > 0$, and we need to show the existences of δ_l and δ_h that make (1) and (2) hold.

When H is the actual world, comparing with σ'_μ , the probability that each contingent agent votes for **A** is increased by $P_{hH} \cdot \delta_h - P_{lH} \cdot \delta_l$ in σ' . Thus, the total expected fraction of votes for **A** is increased by

$$\alpha_C \cdot (P_{hH} \cdot \delta_h - P_{lH} \cdot \delta_l).$$

Similarly, when L is the actual world, similar calculations reveal that the total expected fraction of votes for **A** is increased by

$$\alpha_C \cdot (P_{hL} \cdot \delta_h - P_{lL} \cdot \delta_l).$$

Since the expected fraction of votes for **A** is exactly μ for σ'_μ , we need to choose δ_h and δ_l such that

$$\begin{cases} \alpha_C \cdot (P_{hH} \cdot \delta_h - P_{lH} \cdot \delta_l) > 0 \\ \alpha_C \cdot (P_{hL} \cdot \delta_h - P_{lL} \cdot \delta_l) < 0 \end{cases}.$$

This can always be done due to the positive correlation $P_{hH} > P_{hL}$ and $P_{lH} < P_{lL}$. In particular, if we set $\delta_h = \delta_l \cdot \frac{P_{lH}}{P_{hH}}$, the first inequality would become equality, while the second inequality holds due to the positive correlation. By slightly increasing δ_h , we can make both inequalities hold. During these adjustments, we just need to make sure the two constants δ_h and δ_l are small enough such that β_h and β_l are valid probabilities. Full proof is available in the full version.

EXAMPLE 6. *In this example, we follow the setting of case (2) in Example 5 to illustrate the construction of the strategy σ' . Recall that $\alpha_F = 0.2$, $\alpha_U = 0.3$, and $\alpha_C = 0.5$. The signal distribution $P_{hH} = 0.75$ and $P_{lL} = 0.8$. The threshold $\mu = 0.6$.*

In the first step, let $\sigma'_\mu = (0.8, 0.8)$. We could verify that $\alpha_F + \alpha_C \cdot 0.8 = 0.2 + 0.5 \times 0.8 = 0.6 = \mu$.

In the second step, let $\delta_l = 0.3$. Then $\delta_h = \delta_l \cdot \frac{P_{lH}}{P_{hH}} = 0.1$. Then $\sigma' = (0.5, 0.9)$. Then we have

$$P_{hH} \cdot \delta_h - P_{lH} \cdot \delta_l = 0.75 \times 0.1 - 0.25 \times 0.3 = 0.$$

$$P_{hL} \cdot \delta_h - P_{lL} \cdot \delta_l = 0.2 \times 0.1 - 0.8 \times 0.3 = -0.22 < 0.$$

Finally, we increase δ_h by 0.06. Then $\delta_l = 0.3$, $\delta_h = 0.16$, and $\sigma' = (0.5, 0.96)$. We have

$$P_{hH} \cdot \delta_h - P_{lH} \cdot \delta_l = 0.75 \times 0.16 - 0.25 \times 0.3 = 0.05 > 0.$$

$$P_{hL} \cdot \delta_h - P_{lL} \cdot \delta_l = 0.2 \times 0.16 - 0.8 \times 0.3 = -0.208 < 0.$$

Therefore, $\sigma' = (0.5, 0.96)$ satisfies the condition.

4 PROBABILITY ANALYSIS ON FIDELITY

In this section, we analyze the condition that a strategy profile is of high fidelity and apply the analysis to the most common forms of non-strategic voting: informative voting and sincere voting. We show that neither informative nor sincere voting can lead to the informed majority decision

in every instance, and we characterize the conditions where they lead to the informed majority decision. Our results give merit to strategic voting.

In order to characterize the fidelity, we introduce the notion of the *excess expected vote share*. Given a world state k , the excess expected vote share is the expected vote share the informed majority decision alternative attracts under state k minus the threshold of the alternative.

Definition 4 (Excess expected vote share). Given an instance of N agents, and a strategy profile Σ , let random variable X_n^N be "agent n votes for A": $X_n^N = 1$ if agent n votes for A, and $X_n^N = 0$ if n votes for R. Then the excess expected vote share is defined as follows:

$$f_H^N = \frac{1}{N} \sum_{n=1}^N E[X_n^N \mid H] - \mu. \quad (1)$$

$$f_L^N = \frac{1}{N} \sum_{n=1}^N E[1 - X_n^N \mid L] - (1 - \mu). \quad (2)$$

Specifically, f_H^N is the excess expected vote share of A condition on world state H , and f_L^N is the excess expected vote share of R condition on world state L . For technical convenience, we define $f^N = \min(f_H^N, f_L^N)$.

Our next result shows that we can judge whether the fidelity of a strategy profile sequence converges to 1 with the tendency of its excess expected vote share (or more precisely, the lower limit of $\sqrt{N} \cdot f^N$). If $\sqrt{N} \cdot f^N$ has a lower limit of $+\infty$, Then the fidelity of the profiles in the sequence converges to 1. Otherwise, the fidelity is likely not to converge to 1.

THEOREM 4. *Given an arbitrary sequence of instances and arbitrary sequence of strategy profiles $\{\Sigma_N\}_{N=1}^\infty$, let f^N be the excess expected vote share for each Σ_N .*

- *If $\liminf_{N \rightarrow \infty} \sqrt{N} \cdot f^N = +\infty$, the fidelity of Σ_N converges to 1, i.e., $\lim_{N \rightarrow \infty} A(\Sigma_N) = 1$.*
- *If $\liminf_{N \rightarrow \infty} \sqrt{N} \cdot f^N < 0$ (including $-\infty$), $A(\Sigma_N)$ does NOT converge to 1.*
- *If $\liminf_{N \rightarrow \infty} \sqrt{N} \cdot f^N \geq 0$ (not including $+\infty$), and the variance of $\sum_{n=1}^N X_n^N$ is at least proportional to N , $A(\Sigma_N)$ does NOT converge to 1.*

REMARK. *Although Theorem 4 does not cover the case when $\liminf_{N \rightarrow \infty} \sqrt{N} \cdot f^N \geq 0$ and the variance of $\sum_{n=1}^N X_n^N$ is not large enough, we argue that this case is very special and rare. In this case, $(\inf f^N)$ converges to 0 at the rate of $O(\frac{1}{\sqrt{N}})$, which means the expected vote share of an alternative is almost equal to the threshold. Moreover, the strategies of the agents have low randomness in total. Therefore, we believe that Theorem 4 covers the most interesting cases of a sequence of strategy profiles.*

PROOF SKETCH. Recall that $A(\Sigma) = P_L \cdot \lambda_L^R(\Sigma) + P_H \cdot \lambda_H^A(\Sigma)$. Note that $\lambda_H^A(\Sigma)$ is the probability that the total vote share on A exceeds the threshold μ when the world state is H . Therefore, we can write $\lambda_H^A(\Sigma)$ using the following formula. $\lambda_L^R(\Sigma)$ can be written using a similar formula.

$$\lambda_H^A(\Sigma_N) = \Pr \left[\sum_{n=1}^N X_n^N \geq \mu \cdot N \mid H \right] = \Pr \left[\sum_{n=1}^N X_n^N - \sum_{n=1}^N E[X_n^N \mid H] \geq -f_H^N \cdot N \mid H \right]$$

For the first and the second case, we apply the Hoeffding Inequality. For the first case, we show that both $\lambda_H^A(\Sigma)$ and $\lambda_L^R(\Sigma)$ are lower bounded by a function of N that converges to 1. For the second case, we show that either $\lambda_H^A(\Sigma)$ and $\lambda_L^R(\Sigma)$ is upper bounded by a constant smaller than 1.

For the third case, we apply the Berry-Esseen Theorem [Berry 1941; Esseen 1942], which bounds the difference between the distribution of the sum of independent random variables and the normal distribution. Therefore, for some constant δ and infinitely many N , $\lambda_H^A(\Sigma_N)$ (or $\lambda_L^R(\Sigma_N)$) will not

deviate from $1 - \Phi(\delta)$ too much and is bounded away from 1 by a constant. Φ is the CDF of the standard normal distribution. The requirements for the variance in the third case are from the Berry-Esseen Theorem.

The full proof of Theorem 4 and the detailed definition of the Berry-Esseen Theorem are available in the full version. $\square$

Theorem 4 provides a criterion for judging whether a strategy profile sequence is of high fidelity. If we apply Theorem 2 to each case of Theorem 4, we directly get a criterion for judging whether a regular strategy profile sequence is an ε -strong equilibrium.

COROLLARY 1. *Given an arbitrary sequence of instances and an arbitrary regular sequence of strategy profiles $\{\Sigma_N\}_{N=1}^\infty$, let f^N be defined for each Σ_N .*

- *If $\liminf_{N \rightarrow \infty} \sqrt{N} \cdot f^N = +\infty$, then for every N , Σ_N is an ε -strong BNE with $\varepsilon = o(1)$.*
- *If $\liminf_{N \rightarrow \infty} \sqrt{N} \cdot f^N < 0$ (including $-\infty$), there are infinitely many N such that Σ_N is NOT an ε -strong BNE with constant ε .*
- *If $\liminf_{N \rightarrow \infty} \sqrt{N} \cdot f^N \geq 0$ (not including $+\infty$), and the variance of $\sum_{n=1}^N X_n^N$ is at least proportional to N , there are infinitely many N such that Σ_N is NOT an ε -strong BNE with constant ε .*

EXAMPLE 7. *In this example, we use Theorem 4 to bound the fidelity of different cases in Example 5. Note that in each N the ratio of friendly, unfriendly, and contingent agents is fixed to be 2 : 3 : 5, and agents of the same type play the same strategy for different N . Therefore, the excess expected vote share of profile Σ_N and Σ'_N is independent of N .*

*Case 1: $P_{IH} = 0.1, P_{hH} = 0.9$. In both world states, we have $E[X_n^N | H] = 1$ for friendly agents and $E[X_n^N | H] = 0$ for unfriendly agents. Contingent agents vote for **A** with probability $P_{hH} = 0.9$ in H state and for **R** with probability $P_{IL} = 0.8$ in L state. Therefore, $f^N > 0$, and $\sqrt{N} \cdot f^N$ goes to $+\infty$.*

$$f_H^N = 0.2 + 0.5 \times 0.9 - 0.6 = 0.05 \quad f_L^N = 0.3 + 0.5 \times 0.8 - 0.4 = 0.3.$$

Case 2: $P_{IH} = 0.25, P_{hH} = 0.75$. For Σ_N , we have $f^N < 0$, and $\sqrt{N} \cdot f^N$ goes to $-\infty$.

$$f_H^N = 0.2 + 0.5 \times 0.75 - 0.6 = -0.025 \quad f_L^N = 0.3 + 0.5 \times 0.8 - 0.4 = 0.3.$$

And for the deviating strategy profile Σ'_N , we have $f^N > 0$, and $\sqrt{N} \cdot f^N$ goes to $+\infty$.

$$\begin{aligned} f_H^N &= 0.2 + 0.5 \cdot (0.75 \times 0.96 + 0.25 \times 0.48) - 0.6 = 0.02 \\ f_L^N &= 0.3 + 0.5 \cdot (0.8 \times 0.52 + 0.2 \times 0.04) - 0.4 = 0.112. \end{aligned}$$

In Case 1, the regular strategy profile Σ_N lies in the first case of Theorem 4, has an fidelity converging to 1, and is an ε -strong BNE with $\varepsilon = o(1)$. In Case 2, Σ_N lies in the second case of Theorem 4 and is dominated by the deviating strategy profile Σ'_N . This is in accordance with our observation in Figure 1 and Example 5.

Although Theorem 4 (and Corollary 1) do not cover all the strategy profile sequences, the following result provides a dichotomy for symmetric profile sequences to judge fidelity. Given a symmetric strategy profile Σ_N induced by strategy $\sigma = (\beta_l, \beta_h)$, we can compute the excess expected vote share of Σ_N . Recall the definition of excess expected vote share:

$$f_H^N = \frac{1}{N} \sum_{n=1}^N E[X_n^N | H] - \mu.$$

In H state, an agent with signal h votes for **A** with probability β_h , and an agent with signal l votes for **A** with probability β_l . Therefore, $E[X_n^N | H] = P_{hH} \cdot \beta_h + P_{lH} \cdot \beta_l$. Then we have

$$f_H^N = \frac{1}{N} \sum_{n=1}^N (P_{hH} \cdot \beta_h + P_{lH} \cdot \beta_l) - \mu = P_{hH} \cdot \beta_h + P_{lH} \cdot \beta_l - \mu.$$

With similar reasoning, we can compute f_L^N and f^N :

$$f_L^N = (P_{hL} \cdot (1 - \beta_h) + P_{lL} \cdot (1 - \beta_l)) - (1 - \mu), \quad f^N = \min(f_H^N, f_L^N).$$

An interesting observation for the symmetric strategy profiles is that its excess expected vote share is independent of the number of agents N . This is because when every agent plays the same strategy, the expectation of every X_n^N is the same. For simplicity, given a sequence of symmetric strategy profiles $\{\Sigma_N\}$, we write its excess expected vote share as f_H, f_L , and f .

COROLLARY 2. *For an arbitrary strategy σ and an arbitrary sequence of instances, let $\{\Sigma_N\}$ be the sequence of symmetric strategy profile Σ_N induced by σ , and f be the excess expected vote share of $\{\Sigma_N\}$.*

- If $f > 0$, $A(\Sigma_N)$ converges to 1.
- If $f \leq 0$, $A(\Sigma_N)$ does not converge to 1.

PROOF SKETCH. The proof of Corollary 2 works by showing that each case of a symmetric strategy profile falls into some case of Theorem 4. When $f > 0$, $\sqrt{N} \cdot f^N \rightarrow +\infty$. When $f < 0$, $\sqrt{N} \cdot f^N \rightarrow -\infty$. And when $f = 0$, $\sqrt{N} \cdot f^N = 0$. The variance requirement is also satisfied. (Otherwise, the strategy σ must be always voting for the same candidate. This directly implies $f < 0$, which is a contradiction.) Full proof is available in the full version. $\square$

4.1 Case Study: Informative Voting and Sincere Voting

In this section, we study the two most common non-strategic voting schemes – informative voting and sincere voting under our information structure. We show that both voting schemes lead to the informed majority decision if and only if certain conditions are satisfied.

In informative voting, all agents play the strategy $\sigma = (0, 1)$. When the world state is H , an agent receives signal h and votes for **A** with probability P_{hH} . Therefore, the excess expected vote share in the H state is $f_H = P_{hH} - \mu$. Similarly, the excess expected vote share in the L state is $f_L = P_{lL} - (1 - \mu) = \mu - P_{hL}$. Applying Corollary 2, we get the following statement.

COROLLARY 3. *For an arbitrary sequence of instances, let $\{\Sigma_N\}$ be the sequence of informative voting profile. Then, the fidelity $A(\Sigma_N)$ converges to 1 if and only if $P_{hH} > \mu > P_{hL}$.*




Fig. 2. Illustration of Corollary 3.

Corollary 3 forms a comparison with Theorem 2. Strategic behavior always leads to the informed majority decision, while non-strategic informative voting achieves the same only when the majority vote threshold is “unbiased” compared with the signal distribution.

In sincere voting, an agent votes as if she is making the decision individually. A sincere agent chooses the alternative that maximizes the expected utility conditioned on the signal. The expected utility of an agent making an individual decision conditioned on signal m is

$$\begin{aligned} u_n(\mathbf{A} \mid m) &= \Pr[W = L \mid m] \cdot v_n(\mathbf{A}, L) + \Pr[W = H \mid m] \cdot v_n(\mathbf{A}, H). \\ u_n(\mathbf{R} \mid m) &= \Pr[W = L \mid m] \cdot v_n(\mathbf{R}, L) + \Pr[W = H \mid m] \cdot v_n(\mathbf{R}, H). \end{aligned}$$

Definition 5. A strategy profile Σ is *sincere* if for any agent i , conditioned that i receives signal m , i votes for $\mathbf{A}$ if $u_n(\mathbf{A} \mid m) > u_n(\mathbf{R} \mid m)$ and votes for $\mathbf{R}$ otherwise.

A sincere strategy profile is not always symmetric, because the sincere behavior of agents not only depends on his/her signal but also on his/her utility. Therefore, sincere agents with different utility functions may play different strategies. Given the assumption of $P_{hH} > P_{hL}$, we have $\Pr[W = L \mid l] > \Pr[W = L \mid h]$ and $\Pr[W = H \mid l] < \Pr[W = H \mid h]$. As a result, $u_n(\mathbf{A} \mid l) < u_n(\mathbf{A} \mid h)$ and $u_n(\mathbf{R} \mid l) > u_n(\mathbf{R} \mid h)$. Therefore, a sincere voter would play one of the five strategies below based on her utility function v_n .

- (1) If $u_n(\mathbf{A} \mid l) < u_n(\mathbf{R} \mid l)$, and $u_n(\mathbf{A} \mid h) < u_n(\mathbf{R} \mid h)$, an agent always votes for $\mathbf{R}$.
- (2) If $u_n(\mathbf{A} \mid l) < u_n(\mathbf{R} \mid l)$, and $u_n(\mathbf{A} \mid h) = u_n(\mathbf{R} \mid h)$, an agent votes for $\mathbf{R}$ under signal l and votes arbitrarily under signal h .
- (3) If $u_n(\mathbf{A} \mid l) < u_n(\mathbf{R} \mid l)$, and $u_n(\mathbf{A} \mid h) > u_n(\mathbf{R} \mid h)$, an agent votes informatively.
- (4) If $u_n(\mathbf{A} \mid l) = u_n(\mathbf{R} \mid l)$, and $u_n(\mathbf{A} \mid h) > u_n(\mathbf{R} \mid h)$, an agent votes arbitrarily under signal l , and vote for $\mathbf{A}$ under signal h .
- (5) If $u_n(\mathbf{A} \mid l) > u_n(\mathbf{R} \mid l)$, and $u_n(\mathbf{A} \mid h) > u_n(\mathbf{R} \mid h)$, an agent always votes for $\mathbf{A}$.

A sincere profile is also a regular profile, as friendly agents always vote for $\mathbf{A}$, and unfriendly agents always vote for $\mathbf{R}$ in their individual decisions. Therefore, applying Theorem 2, we have the following statement.

COROLLARY 4. For an arbitrary sequence of instances, let $\{\Sigma_N\}$ be the sequence of sincere strategy profiles. Then the fidelity $A(\Sigma_N)$ converges to 1 if and only if Σ_N is an ε -strong Bayes Nash with $\varepsilon = o(1)$.

Corollary 4 tells us that sincere voting performs as well as strategic voting if and only if itself is also strategic. The following example illustrates different behaviors of sincere voters by “manipulating” their utility functions under the same world state and signal distribution and gives examples where sincere voting succeeds and fails.

EXAMPLE 8. Consider the following scenario. The world state prior $P_H = P_L = 0.5$. The signal distribution $P_{hH} = P_{lL} = 0.8$, and $P_{lH} = P_{hL} = 0.2$. By the Bayes Theorem, we compute the probability of a world state conditioned on a private signal as follows.

$$\begin{aligned} \Pr[W = L \mid l] &= 0.8, & \Pr[W = H \mid l] &= 0.2 \\ \Pr[W = L \mid h] &= 0.2, & \Pr[W = H \mid h] &= 0.8. \end{aligned}$$

We assume all agents are contingent and share the same utility function, and consider three different cases as shown in Table 7. Suppose Σ is a strategy profile where all agents vote sincerely.

(Winner, World State)	(A, H)	(A, L)	(R, H)	(R, L)
Case 1	1	0	0	1
Case 2	5	1	0	2
Case 3	4	1	0	2

Table 7. Utility functions for three cases.

In case 1, $u_n(\mathbf{A} \mid h) = u_n(\mathbf{R} \mid l) = 0.8$, and $u_n(\mathbf{A} \mid l) = u_n(\mathbf{R} \mid h) = 0.2$. Therefore, every sincere voter votes informatively, and Σ is also informative voting. By corollary 3, Σ leads to the informed majority decision if and only if the threshold $0.2 < \mu < 0.8$.

In case 2, $u_n(\mathbf{A} \mid l) = 1.8$, $u_n(\mathbf{R} \mid l) = 1.6$, $u_n(\mathbf{A} \mid h) = 4.2$, and $u_n(\mathbf{R} \mid h) = 0.4$. Therefore, all sincere agents always vote for $\mathbf{A}$ even if they are contingent. In this case, $\mathbf{A}$ is always the winner, and $A(\Sigma)$ does not converge to 1.

In case 3, $u_n(\mathbf{A} \mid l) = u_n(\mathbf{R} \mid l) = 1.6$, $u_n(\mathbf{A} \mid h) = 3.4$, and $u_n(\mathbf{R} \mid h) = 0.4$. In this case, a sincere agent votes $\mathbf{A}$ under signal h , and votes arbitrarily under signal l . Then for any $0.2 < \mu < 1$, Σ induced by strategy $\sigma = (\beta_l, 1)$ leads to the informed majority decision with high probability, where $\beta_l \geq 0$ satisfies $\beta_l > 5\mu - 4$ and $\beta_l < 1.25\mu - 0.25$. These conditions guarantee the excess expected vote share of Σ to be strictly positive.

As shown in Example 8, sincere agents can have drastically different behaviors in different scenarios. Nevertheless, once we know the strategy of each agent, we can apply Theorem 4 to analyze the probability of a sequence of sincere voting leading to the informed majority decision.

5 NON-BINARY WORLD STATES AND NON-BINARY SIGNALS

In this section, we discuss how we extend our model and results to a setting with non-binary world states and non-binary signals. We follow the setting of Schoenebeck and Tao [2021]. The largest difference in the non-binary setting is that the preferences of agents form a spectrum along multiple world states. Different agents have different thresholds of world states in which they switch the preferred alternatives.

World State. There are K possible world states $\mathcal{W} = \{1, 2, \dots, K\}$. The higher the world state is, the more $\mathbf{A}$ is preferred to $\mathbf{R}$. We use k to denote a generic world state. Let $P_k = \Pr[W = k]$ be the common prior of the world state. We assume $P_k > 0$ for every $k \in \mathcal{W}$.

Signal. Every agent receives a signal from $\mathcal{S} = \{1, 2, \dots, M\}$. We use m to denote a generic signal. Signals are i.i.d conditioned on the world state. Let $P_{mk} = \Pr[S_n = m \mid W = k]$ be the probability that an agent receives signal m given world state is k . The assumption of $P_{hH} > P_{hL}$ in the binary state is extended to the following assumption of stochastic dominance, which requires signals to be positively correlated to world states.

ASSUMPTION 1 (STOCHASTIC DOMINANCE). For any agent n and any world states $k_1 > k_2$,

$$\Pr[S_n \geq m \mid W = k_1] > \Pr[S_n \geq m \mid W = k_2].$$

Utility. Every agent n has a utility function $v_n : \mathcal{W} \times \{\mathbf{A}, \mathbf{R}\} \rightarrow \{0, 1, \dots, B\}$, where B is the positive integer upper bound. We assume that $\mathbf{A}$ is more preferable in a higher world state than in a lower state, and $\mathbf{R}$ is the opposite: for any k_1 and k_2 with $k_1 > k_2$, $v_n(k_1, \mathbf{A}) > v_n(k_2, \mathbf{A})$, and $v_n(k_1, \mathbf{R}) < v_n(k_2, \mathbf{R})$. We also assume $v_n(k, \mathbf{A}) \neq v_n(k, \mathbf{R})$ for any k and any n .

Fraction of agents. Given a world state k , let $\alpha_k^{\mathbf{A}}$ and $\alpha_k^{\mathbf{R}}$ be the *approximated* fraction of agents prefer $\mathbf{A}$ ($\mathbf{R}$, respectively) in world state k . We assume $\alpha_k^{\mathbf{A}}$ and $\alpha_k^{\mathbf{R}}$ are independent from N . Formally,

given N , let $N(k, \mathbf{A}) = \{n \mid v_n(k, \mathbf{A}) > v_n(k, \mathbf{R})\}$ be the set of agents preferring $\mathbf{A}$ in k ($N(k, \mathbf{R})$ defined similarly). We have $|N(k, \mathbf{R})| = \lfloor \alpha_k^{\mathbf{R}} \cdot N \rfloor$ and $|N(k, \mathbf{A})| = N - |N(k, \mathbf{R})|$. Naturally, $\alpha_k^{\mathbf{A}}$ increases, and $\alpha_k^{\mathbf{R}}$ decreases as k increases. We assume that $\alpha_k^{\mathbf{A}}$ and $\alpha_k^{\mathbf{R}}$ are common knowledge.

Majority vote and informed majority decision. We study the majority vote with threshold μ . If at least $\mu \cdot N$ agents vote for $\mathbf{A}$, $\mathbf{A}$ is announced to be the winner; otherwise, $\mathbf{R}$ is announced to be the winner. The informed majority decision is defined on each world state k . Given a world state k , if $\alpha_k^{\mathbf{A}} > \mu$, we say $\mathbf{A}$ is the informed majority decision; otherwise, we say $\mathbf{R}$ is the informed majority decision. We assume that $\alpha_k^{\mathbf{A}} \neq \mu$ for all $k \in \mathcal{W}$, and the rounding between $\alpha_k^{\mathbf{A}}$ and $|N(k, \mathbf{A})|$ ($\alpha_k^{\mathbf{R}}$ and $|N(k, \mathbf{R})|$, respectively) does not flip the informed majority decision.

Types of agents. Let $\mathcal{L} = \{k \in \mathcal{W} \mid \alpha_k^{\mathbf{A}} < \mu\}$ and $\mathcal{H} = \{k \in \mathcal{W} \mid \alpha_k^{\mathbf{A}} > \mu\}$ be the sets of world states where $\mathbf{R}$ ($\mathbf{A}$, respectively) is the informed majority decision. We only consider the case where both $\mathcal{L}$ and $\mathcal{H}$ are non-empty. (Otherwise, an alternative is unanimously the informed majority decision, and there is no uncertainty.) There is a threshold partitioning $\mathcal{W}$ into two sets. Let $L = \max\{k \in \mathcal{L}\}$ to be the largest world where $\mathbf{R}$ is the informed majority decision, and $H = \min\{k \in \mathcal{H}\}$ to be the smallest world where $\mathbf{A}$ is the informed majority decision. We have $H = L + 1$.

Similarly, for an agent n , let $\mathcal{L}_n = \{k \in \mathcal{W} \mid v_n(k, \mathbf{R}) > v_n(k, \mathbf{A})\}$ and $\mathcal{H}_n = \{k \in \mathcal{W} \mid v_n(k, \mathbf{R}) < v_n(k, \mathbf{A})\}$. $\mathcal{L}_n$ ($\mathcal{H}_n$, respectively) is the set of world states where $\mathbf{R}$ ($\mathbf{A}$, respectively) is preferred by n . For an agent n , let $L_n = \max\{k \in \mathcal{L}_n\}$ be the largest world where n prefers $\mathbf{R}$, and $H_n = \min\{k \in \mathcal{H}_n\}$ to be the smallest world where n prefers $\mathbf{A}$. Specifically, let $L_n = 0$ if $\mathcal{L}_n = \emptyset$ and $H_n = K + 1$ if $\mathcal{H}_n = \emptyset$. We have $H_n = L_n + 1$.

- (1) We say an agent n is *(candidate) friendly* if $\mathcal{L} \cap \mathcal{H}_n \neq \emptyset$. This says that there exists a world state k where n prefer $\mathbf{A}$ while the informed majority decision is $\mathbf{R}$.
- (2) Similarly, an agent n is *(candidate) unfriendly* if $\mathcal{H} \cap \mathcal{L}_n \neq \emptyset$. This says that there exists a world state k where n prefers $\mathbf{R}$ while the informed majority decision is $\mathbf{A}$.
- (3) Finally, an agent n is *contingent* if $\mathcal{L}_n = \mathcal{L}$. or equivalently, $\mathcal{H}_n = \mathcal{H}$.

Unlike the binary setting, friendly/unfriendly agents in the non-binary setting do not always prefer one alternative. They just have thresholds above or below the majority.

EXAMPLE 9. We extend the COVID policy-making scenario in Example 4 to the non-binary setting. $N = 20$ voters decide whether to accept or reject the more-restrictive policy. The world states $\mathcal{W} = \{1, 2, 3\}$ describe the risk level of the virus, whereas a larger state represents a higher risk. Suppose $P_1 = P_2 = 0.3$, and $P_3 = 0.4$.

Every voter receives a private signal from $\mathcal{S} = \{1, 2, 3, 4\}$. The larger the signal is, the higher the risk is likely to be. Table 8 is the signal distribution given the world state.

World State	Signal 1	Signal 2	Signal 3	Signal 4
1	0.6	0.2	0.1	0.1
2	0.4	0.2	0.2	0.2
3	0.1	0.2	0.3	0.4

Table 8. Signal distribution.

The majority threshold is $\mu = 0.6$. Therefore, $\mathbf{A}$ is the informed majority decision if and only if at least 12 agents prefer $\mathbf{A}$ to $\mathbf{R}$. The voters are categorized into four different groups. Each group has five voters, and voters in the same group share the same utility shown in Table 9. The larger the group index is, the more voters prefer $\mathbf{A}$ to $\mathbf{R}$.

Table 10 shows the preferences of each group and the informed majority decision under each world state. The preference of each group comes from the comparison of utilities in Table 9. The informed majority decision is the aggregation of group preferences. Since each group has five voters, and the majority threshold is 12 agents, **A** needs to be preferred by at least three groups to become the informed majority decision. Therefore, **A** is the informed majority decision only in state 3. By comparing the preferences of each group and the informed majority decision, we know that group 1 voters are unfriendly agents, group 2 voters are contingent agents, and group 3 and 4 voters are friendly agents.

Winner	A			R		
World State	1	2	3	1	2	3
Group 1	1	2	3	8	6	4
Group 2	2	3	4	6	4	2
Group 3	2	5	8	4	3	2
Group 4	4	6	9	3	2	1

Table 9. Utility function of each group

World State	1	2	3
Group 1	R	R	R
Group 2	R	R	A
Group 3	R	A	A
Group 4	A	A	A
Informed Majority	R	R	A

Table 10. Preference of each group and the majority.

Strategy. In the non-binary setting, a strategy can be represented as a vector $\sigma = (\beta_1, \beta_2, \dots, \beta_M)$, where β_m is the probability that the agent votes for **A** when receiving signal m . A strategy profile is the vector of strategies of all agents. $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$.

The definition of the regular strategy profile remains the same as in the binary setting: friendly agents always vote for **A**, and unfriendly agents always vote for **R**.

Fidelity and Expected Utility. Given a strategy profile Σ , let $\lambda_k^{\mathbf{A}}(\Sigma)$ ($\lambda_k^{\mathbf{R}}(\Sigma)$, respectively) be the probability that **A** (**R**, respectively) becomes the winner when world state is k . We can define the fidelity and the expected utility in the same manner as in the binary setting.

$$A(\Sigma) = \sum_{k \in \mathcal{L}} P_k \cdot \lambda_k^{\mathbf{R}}(\Sigma) + \sum_{k \in \mathcal{H}} P_k \cdot \lambda_k^{\mathbf{A}}(\Sigma).$$

$$u_n(\Sigma) = \sum_{k=1}^K P_k (\lambda_k^{\mathbf{A}}(\Sigma) \cdot v_n(k, \mathbf{A}) + \lambda_k^{\mathbf{R}}(\Sigma) \cdot v_n(k, \mathbf{R})).$$

Excess Expected Vote Share. Similarly, the excess expected vote share is the expected vote share that an alternative attracts under state k minus the threshold of the alternative. In different instances, the informed majority decision may change in different world states. Therefore, we define the excess expected vote share for both **A** and **R** in every world state.

$$f_{k\mathbf{A}}^N = \frac{1}{N} \sum_{n=1}^N E[X_n^N | k] - \mu. \quad f_{k\mathbf{R}}^N = \frac{1}{N} \sum_{n=1}^N E[1 - X_n^N | k] - (1 - \mu).$$

For world states $k \in \mathcal{H}$ where **A** is the informed majority decision, we care about $f_{k\mathbf{A}}^N$; and for states $k \in \mathcal{L}$, we care about $f_{k\mathbf{R}}^N$. Therefore, we define f^N to be the smallest excess expected vote share among those we care about, i.e. $f^N = \min \left(\min_{k \in \mathcal{H}} (f_{k\mathbf{A}}^N), \min_{k \in \mathcal{L}} (f_{k\mathbf{R}}^N) \right)$. We use $f_{k\mathbf{A}}, f_{k\mathbf{R}}$, and f to denote excess expected vote share in symmetric profile sequences.

Instance and Sequence of Strategy Profiles. Let $\{\mathcal{I}_N\}_{N=1}^\infty$ (or $\{\mathcal{I}_N\}$ for short) be a sequence of instances, where each $\mathcal{I}_N$ is an instance of N agents. The instances in a sequence share the same majority threshold μ , world state prior distribution $\{P_k\}$, signal prior distribution $\{P_{mk}\}$, and approximated type fractions $(\alpha_k^{\mathbf{A}}, \alpha_k^{\mathbf{R}})$. Same to the binary setting, we do not regard agents in

different instances as related and have no additional assumption on the utility functions of agents. We define a sequence of strategy profile $\{\Sigma_N\}_{N=1}^{\infty}$ based on the instance sequence, where for each N , Σ_N is a strategy profile in $\mathcal{I}_N$.

Our positive results on strategic voting can be extended to the non-binary setting. Theorem 5 states the equivalence of fidelity converging to 1 and the ε -strong Bayes Nash with $\varepsilon = o(1)$. Theorem 6 guarantees the existence of the regular profile sequence with fidelity converging to 1.

THEOREM 5. *Given an arbitrary sequence of instance and an arbitrary regular strategy profile sequence $\{\Sigma_N\}_{N=1}^{\infty}$, let $\{A(\Sigma_N)\}_{N=1}^{\infty}$ be the sequence of the fidelities of Σ_N .*

- *If $\lim_{N \rightarrow \infty} A(\Sigma_N) = 1$, then for every N , Σ_N is an ε -strong BNE with $\varepsilon = o(1)$.*
- *If $\lim_{N \rightarrow \infty} A(\Sigma_N) = 1$ does not hold, then there exist infinitely many N such that Σ_N is NOT an ε -strong BNE for some constant ε .*

THEOREM 6. *Given any arbitrary sequence of instances, there always exists a sequence of regular strategy profiles $\{\Sigma'_N\}_{N=1}^{\infty}$ such that $A(\Sigma'_N)$ converges to 1.*

Theorem 5 and Theorem 6 preserves the same reasoning as Theorem 2 and Theorem 3. More details as well as the characterization of the fidelity in the non-binary setting are available in the full version of this paper.

6 CONCLUSION AND FUTURE WORK

We study the binary voting game where agents can coordinate in groups. We show that strategic voting always leads to the informed majority decision, while non-strategic behaviors sometimes fail. In particular, we show that a strategy profile is an ε -strong Bayes Nash Equilibrium with small ε if and only if it leads to the “correct” decision with high probability. Moreover, we analyze the fidelity of the strategy profile and provide criteria for judging whether a strategy profile is an equilibrium based on excess expected vote share. Applying the analysis to non-strategic voting, we characterize the conditions that informative and sincere voting lead to the informed majority decision. Our results stand on the framework where agents have endogenous preferences over outcomes contingent on some underlying state.

One limitation of our work is that our results are restricted to a setting with two alternatives. An interesting yet challenging future direction is to study the impact of strategic behavior in a setting with more than two alternatives. We expect more complicated results as Goertz and Maniquet [2014] show that informative voting may be an equilibrium but leads to the wrong alternative in a model of three alternatives.

Another interesting direction is to explore strategic iterative voting with information uncertainty. We expect iterative voting to be more powerful in aggregating information and able to simulate some sophisticated mechanisms. For example, the mechanism in Schoenebeck and Tao [2021] can be regarded as a two-round voting game where only the second round counts, and every agent votes informatively in the first round and plays a surprisingly popular strategy in the second round. Kavner and Xia [2021] show a surprising result that strategic behaviors increase the social welfare of agents in iterative voting on average. Nevertheless, the behavior of strategic agents is even more complicated in iterative voting, and the analysis of equilibria will be highly challenging.

ACKNOWLEDGMENTS

We thank all anonymous reviewers for their helpful comments and suggestions. GS acknowledges NSF #2007256 and NSF #1452915 for support. LX acknowledges NSF #1453542, #2007994, and #2106983, and a Google Research Award for support. The research of Biaoshuai Tao was supported by the National Natural Science Foundation of China (Grant No. 62102252).

REFERENCES

- Kenneth J. Arrow. 2012. *Social Choice and Individual Values, Third Edition*. Yale University Press.
- David Austen-Smith and Jeffrey S. Banks. 1996. Information Aggregation, Rationality, and the Condorcet Jury Theorem. *The American Political Science Review* 90, 1 (1996), 34–45.
- Salvador Barberà, Hugo Sonnenschein, and Lin Zhou. 1991. Voting by Committees. *Econometrica* 59, 3 (1991), 595–609.
- Ruth Ben-Yashar and Mor Zahavi. 2011. The Condorcet jury theorem and extension of the franchise with rationally ignorant voters. *Public Choice* 148, 3 (2011), 435–443.
- Andrew C. Berry. 1941. The Accuracy of the Gaussian Approximation to the Sum of Independent Variates. *Trans. Amer. Math. Soc.* 49, 1 (1941), 122–136.
- J. Eric Bickel. 2007. Some comparisons among quadratic, spherical, and logarithmic scoring rules. *Decision Analysis* 4, 2 (2007), 49–65.
- Philip J. Boland, Frank Proschan, and Y. L. Tong. 1989. Modelling Dependence in Simple and Indirect Majority Systems. *Journal of Applied Probability* 26, 1 (1989), 81–88.
- Marquis de Condorcet. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris: L'Imprimerie Royale.
- John Duggan and Cesar Martinelli. 2001. A Bayesian Model of Voting in Juries. *Games and Economic Behavior* 37, 2 (2001), 259–294.
- Carl-Gustav Esseen. 1942. On the Liapunoff limit of error in the theory of probability. *Arkiv för Matematik, Astronomi och Fysik* A28, 9 (1942), 1–19.
- Timothy Feddersen and Wolfgang Pesendorfer. 1997. Voting Behavior and Information Aggregation in Elections With Private Information. *Econometrica* 65, 5 (1997), 1029–1058.
- Timothy Feddersen and Wolfgang Pesendorfer. 1998. Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting. *American Political Science Review* 92, 1 (1998), 23–35.
- Tilmann Gneiting and Adrian E Raftery. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Amer. Statist. Assoc.* 102, 477 (2007), 359–378. <https://doi.org/10.1198/016214506000001437>
- Johanna M.M. Goertz and François Maniquet. 2014. Condorcet Jury Theorem: An example in which informative voting is rational but leads to inefficient information aggregation. *Economics Letters* 125, 1 (2014), 25–28.
- Mark Gradstein and Shmuel Nitzan. 1987. Organizational decision-making quality and the severity of the free-riding problem. *Economics Letters* 23, 4 (1987), 335–339.
- Bernard Grofman, Guillermo Owen, and Scott L. Feld. 1983. Thirteen theorems in search of the truth. *Theory and Decision* 15, 3 (1983), 261–278.
- Hadi Hosseini, Debmalaya Mandal, Nisarg Shah, and Kevin Shi. 2021. Surprisingly Popular Voting Recovers Rankings, Surprisingly!. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 245–251. <https://doi.org/10.24963/ijcai.2021/35>
- Serguei Kaniovski. 2010. Aggregation of correlated votes and Condorcet's Jury Theorem. *Theory and Decision* 69, 3 (2010), 453–468.
- Joshua Kavner and Lirong Xia. 2021. Strategic Behavior is Bliss: Iterative Voting Improves Social Welfare. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 19021–19032.
- Adam Meirowitz. 2002. Informative voting and Condorcet jury theorems with a continuum of types. *Social Choice and Welfare* 19, 1 (2002), 219–236.
- Mark S. Miller and K. Eric Drexler. 1988. Markets and Computation: Agoric Open Systems. In *The Ecology of Computation*, B. A. Huberman (Ed.). North-Holland.
- Nolan Miller, Paul Resnick, and Richard Zeckhauser. 2005. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science* 51, 9 (2005), 1359–1373.
- Nicholas R. Miller. 1986. Information, Electorates, and Democracy: Some Extensions and Interpretations of the Condorcet Jury Theorem. In *Information Pooling and Group Decision Making*, Grofman B. and Owen G. (Eds.). JAI Press, 173–192.
- Roger B. Myerson. 1998. Extended Poisson Games and the Condorcet Jury Theorem. *Games and Economic Behavior* 25, 1 (1998), 111–131.
- Shmuel Nitzan and Jacob Paroush. 1980. Investment in Human Capital and Social Self Protection under Uncertainty. *International Economic Review* 21, 3 (1980), 547–557.
- Shmuel Nitzan and Jacob Paroush. 1984. The significance of independent decisions in uncertain dichotomous choice situations. *Theory and Decision* 17, 1 (1984), 47–60.
- Shmuel Nitzan and Jacob Paroush. 2017. Collective Decision Making and Jury Theorems. In *The Oxford Handbook of Law and Economics: Volume 1: Methodology and Concepts*, Francesco Parisi (Ed.). Oxford University Press.
- Guillermo Owen, Bernard Grofman, and Scott L. Feld. 1989. Proving a distribution-free generalization of the Condorcet Jury Theorem. *Mathematical Social Sciences* 17, 1 (1989), 1–16.

- Charles R. Plott. 1976. Axiomatic Social Choice Theory: An Overview and Interpretation. *American Journal of Political Science* 20, 3 (1976), 511–596.
- Dražen Prelec. 2004. A Bayesian Truth Serum for Subjective Data. *Science* 306, 5695 (2004), 462–466.
- Dražen Prelec, H Sebastian Seung, and John McCoy. 2017. A solution to the single-question crowd wisdom problem. *Nature* 541, 7638 (2017), 532–535.
- Grant Schoenebeck and Biaoshuai Tao. 2021. Wisdom of the Crowd Voting: Truthful Aggregation of Voter Information and Preferences. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 1872–1883.
- Grant Schoenebeck and Fang-Yi Yu. 2021. Learning and Strongly Truthful Multi-Task Peer Prediction: A Variational Approach. In *12th Innovations in Theoretical Computer Science Conference (Leibniz International Proceedings in Informatics, Vol. 185)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 78:1–78:20.
- Grant Schoenebeck, Fang-Yi Yu, and Yichi Zhang. 2021. Information Elicitation from Rowdy Crowds. In *Proceedings of the Web Conference 2021*. Association for Computing Machinery, New York, NY, USA, 3974–3986. <https://doi.org/10.1145/3442381.3449840>
- Lloyd Shapley and Bernard Grofman. 1984. Optimizing group judgmental accuracy in the presence of interdependencies. *Public Choice* 43, 3 (1984), 329–343.
- Jörgen Wit. 1998. Rational Choice and the Condorcet Jury Theorem. *Games and Economic Behavior* 22, 2 (1998), 364–376.
- Jens Witkowski and David C. Parkes. 2012. Peer Prediction Without a Common Prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. Valencia, Spain, 964–981.
- Justin Wolfers and Eric Zitzewitz. 2004. Prediction Markets. *Journal of Economic Perspectives* 18, 2 (June 2004), 107–126. <https://doi.org/10.1257/0895330041371321>
- H. Peyton Young. 1988. Condorcet’s Theory of Voting. *American Political Science Review* 82 (1988), 1231–1244.