

The Art of Two-Round Voting

QISHEN HAN, Rutgers University, United States

GRANT SCHOENEBECK, University of Michigan, United States

BIAOSHUAI TAO, Shanghai Jiao Tong University, China

LIRONG XIA, Rutgers University and DIMACS, United States

We study the voting problem with two alternatives where voters' preferences depend on a not-directly-observable state variable. While equilibria in the one-round voting mechanisms lead to a good decision, they are usually hard to compute and follow. We consider the two-round voting mechanism where the first round serves as a polling stage and the winning alternative only depends on the outcome of the second round. We show that the two-round voting mechanism is a powerful tool for making collective decisions. Firstly, every (approximated) equilibrium in the two-round voting mechanisms (asymptotically) leads to the decision preferred by the majority as if the state of the world were revealed to the voters. Moreover, there exist natural equilibria in the two-round game following intuitive behaviors such as informative voting, sincere voting [Austen-Smith and Banks, 1996], and surprisingly popular strategy [Prelec et al., 2017]. This sharply contrasts the one-round voting mechanisms in the previous literature, where no simple equilibrium is known. Finally, we show that every equilibrium in the standard one-round majority vote mechanism gives an equilibrium in the two-round mechanisms that is not more complicated than the one-round equilibrium. Therefore, the two-round voting mechanism provides a natural equilibrium in every instance including those where one-round voting fails to have a natural solution, and it can reach an informed majority decision whenever one-round voting can. Our experiments on generative AI voters also imply that two-round voting leads to the correct outcome more often than one-round voting under some circumstances.

CONTENTS

Abstract	0
Contents	0
1 Introduction	1
1.1 Our Contribution	2
1.2 Related Work and Discussion	3
2 Preliminaries	4
2.1 One round voting	6
3 Equilibria in One-Round Voting	7
4 Two-round Voting and Equilibria	8
4.1 Equilibria in Two Round Voting	9
4.2 Proof Sketch for Theorem 1	11
4.3 Proof Sketch for Theorem 2 and Theorem 3.	13
5 Two-Round Voting is More Powerful	14
6 General Valuation Utility.	14
7 Experiments	16
7.1 Experimental Setting	16
7.2 Experimental Results	17
7.3 Take-away Messages	17
8 Conclusion and Future Work	18
References	18

1 Introduction

How can people make informed collective decisions? In many scenarios, people’s preferences depend on the information they receive, and a collective decision needs to be made fairly. For example, suppose that a community wants to decide whether or not to adopt a restrictive COVID-19 policy. The preferences of the community members depend on their perceptions of the virality of the pandemic. While the virality is not directly observable, community members receive (noisy) signals about it to form their preferences, and their preferences can still be different even upon receiving the same signal. For example, immunocompromised people may still prefer to adopt the restrictive policy when the signal is “moderately dangerous”, while healthy people may prefer not to adopt the policy after receiving the same signal. Other examples include political elections, recruiting, group activity selection, etc.

A natural idea is to use voting to make collective decisions. When there are two alternatives, majority voting appears to be a natural choice. Majority voting reveals the state of the world (the virality of COVID-19 in the above example) almost surely when the population is large, according to the Condorcet Jury theorem [Condorcet, 1785]. Moreover, when voters are strategic, majority voting can still aggregate voters’ fully informed preferences via strategic behaviors [Feddersen and Pesendorfer, 1997, Han et al., 2023]. In other words, it chooses the majority winner *as if the state of the world were revealed to the voters*. In the literature, such majority decision has been referred to as *full information equivalence* [Feddersen and Pesendorfer, 1997], *majority wish* [Schoenebeck and Tao, 2021], and *informed majority decision* [Han et al., 2023]. In this paper, we will use the term “informed majority decision”.

However, all equilibria of such one-round majority games in previous work are quite complicated and hard to follow. For example, the seminal paper by Austen-Smith and Banks [1996] shows that, under majority voting, the natural and intuitive behavior of *informative voting*, in which every voter votes for his/her signal, and *sincere voting*, in which every voter votes as if he/she is making an individual decision and vote for his/her preferred alternative in expectation, are not Nash equilibria in most situations. The equilibrium constructed by Feddersen and Pesendorfer [1997] requires the computation of a complicated Riemann integral. The equilibrium constructed by Han et al. [2023] involves carefully calculated mixed strategies.

In analyzing strategic behavior, we emphasize the importance of natural and intuitive equilibria. It is quite unlikely that real-world agents can calculate out and play to counter-intuitive mixed or asymmetric strategies, let alone play the same equilibrium. Therefore, these counter-intuitive strategy profiles fail to predict the outcome of strategy agents as they are intended for. Moreover, the simplicity of the mechanism is also a desirable property and a proxy of usability, as the theoretical guarantees of complex mechanisms may not translate into practice.

Can we design *simple mechanisms with natural equilibria to achieve the informed majority decision?*

The definition of a natural equilibrium depends on the scenarios and the purposes. In this paper, we identify the following criteria. Firstly, agents play pure/deterministic strategies. Mixed or randomized strategies are hard for agents to compute and implement and are rarely applied by real-world voters. Secondly, agents play symmetrically, i.e., agents with the same preferences should play the same strategy. Otherwise, the mapping from similar agents to diverse strategies is unclear and unincited, and the equilibrium is unlikely to predict the outcome. Moreover, agents can easily compute the strategies. Two examples of intuitive and natural strategies are informative voting and sincere voting, where the first simply follows the signal, and the second follows the preferences in expectation.

decisions with strategic agents, as every ε -approximate strong Bayes Nash equilibrium, in which no group of agents has incentives to deviate, leads to the informed majority decision with high probability (Theorem 1). The probability converges to 1 as ε goes to 0. Secondly, the two-round voting mechanism has the merit of a simple mechanism in both structure-wise and equilibrium-wise. It has a simple structure that involves simply conducting a poll before the actual vote. Such a polling-voting structure has been widely applied in real-world political elections.

Moreover, we show a natural ε -approximate strong Bayes Nash equilibrium that leads to the informed majority decision (Theorem 2). As the number of voters increases, ε converges to 0. The equilibrium combines the most natural and classical informative voting and sincere voting. In this equilibrium, agents whose preferred alternative changes with the world state vote informatively (vote for their signal) in the first round, and vote sincerely (vote for the alternative they preferred in expectation conditioned on the first-round outcome) in the second round. The “informative+sincere” equilibrium follows the intuition of sharing information and choosing the (expected) preferred alternative. We further discover that the “informative+sincere” equilibrium belongs to a class of “informative + threshold” equilibria (Theorem 3), in which agents vote informatively in the first round and vote for the alternative whose first-round vote exceeds a threshold. The “informative + threshold” class derives more natural equilibria, such as the “informative + surprisingly popular” equilibrium in which agents vote for the “surprisingly popular” alternative in the second round.

Furthermore, we show with mild assumptions that, every one-round equilibrium gives a two-round equilibrium that is not more “complicated” (Theorem 4). Specifically, agents in the two-round voting mechanism can ignore the polling stage and directly play the one-round equilibrium in the second round. Therefore, the two-round voting mechanism provides natural equilibrium in every instance including those where one-round voting fails to have a natural solution, and it can reach an informed majority decision whenever one-round voting can.

We also test the one-round voting and the two-round voting with synthetic experiments where we apply generate AI agents to simulate voters. When the signal distribution is biased, the AI agents powered by the reasoning model reach the informed majority decision more often in the two-round voting than in the one-round voting. We also observe behaviors such as Bayesian update and vote switch (in the second round). Our experiments provide insights into future applications for AI-augmented or proxy votes in the future and serve as references for real-world experiments.

1.2 Related Work and Discussion

The study of the binary voting game with imperfectly informed agents originates from the famous Condorcet Jury Theorem [Condorcet, 1785], which shows the strength of voting when seeking the truth. Condorcet showed that the majority vote can reveal the “correct alternative” with probability converging to one as the number of voters increases, conditioned on (1) agents honestly reflecting their information in the votes (informative voting), and (2) the probability of an agent getting the correct information $p > 0.5$.

The seminal paper by Austen-Smith and Banks [1996] is the first to study this problem in a game-theoretical setting. They show that informative voting is not compatible with the strategic behavior of voters. As a consequence, a large literature focuses on the existence of non-informative equilibrium and its effect on making a good decision. Wit [1998], Feddersen and Pesendorfer [1998], and Myerson [1998] discovers the existence of a mixed strategy equilibrium that reveals the ground truth under different settings. Feddersen and Pesendorfer [1997] propose a unique class of equilibria that reach the informed majority decision. Han et al. [2023] reveal the equivalence between a strategy profile being an equilibrium and reaching a good decision. None of these works guarantee the existence of a “simple” equilibrium.

Multiple works go “beyond” the one-round voting game, investigating and proposing more sophisticated mechanisms to aggregate information. Schoenebeck and Tao [2021] design an information aggregation mechanism that incentivizes informative voting and leads to an informed majority decision with high probability. Their mechanism leverages the “surprisingly popular mechanism” and the median-voter theorem to incentivize informative votes. Coughlan [2000] shows that when agent preferences are restricted to similar, adding an communication phase before voting leads to a good decision. Their result does not extend to scenarios where agents have diverse preferences as in our setting because they adopted a different solution concept (Bayesian-Nash equilibrium) and their techniques are very different from ours. Morgan and Stocken [2008] study a two-round decision-making game in which voters vote in the first-round poll and the individual decision maker chooses the decision in the second round. In their model, agents cannot make a direct decision but only affect the decision-maker by their reported signals. Dekel and Piccione [2000] study the voting game in which agents publicly declare their vote in a predetermined order. They show that a strategy profile is an equilibrium in sequential voting if and only if a corresponding profile is an equilibrium in one-round voting. This implies that sequential voting may not have a simple equilibrium. Both papers are very different from ours in the model and the results. Chopra et al. [2004] studies the necessary knowledge for agents to strategize in the voting problems. Other directions of one-round voting with imperfect information include dependent signals [Kaniowski, 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], voting with more than two alternatives [Goertz and Maniquet, 2014, Young, 1988], and empirical studies [Battaglini et al., 2010, Esponda and Vespa, 2014].

Our paper also aligns with the idea that communication in voting improves aggregation, especially Gerardi and Yariv [2007] which shows that communication before voting helps stabilize the voting outcome. Their results show that equilibrium outcomes in “deliberation+different voting rules” are the same. Goeree and Yariv [2011] further show empirically that these outcomes resemble those in one-round informative voting. They don’t examine whether outcomes lead to good decisions, while we show that all equilibria in the two-round voting lead to good decisions.

Our paper is also related to strategic information transmission where information communication occurs between informed experts and an uninformed decision maker [Austen-Smith, 1993, Battaglini, 2002, 2004, Crawford and Sobel, 1982, Morgan and Stocken, 2008]. Divergent preferences between experts and decision-makers usually lead to information loss. However, even when experts strategically report their information, the decision maker can still collect enough information to make an optimal decision [Baharad et al., 2012, Morgan and Stocken, 2008].

2 Preliminaries

In this section, we introduce the information structure and the one-round voting game model adopted from previous work [Han et al., 2023, Schoenebeck and Tao, 2021].

Alternatives and World States. n agents vote for two alternatives **A** and **R**. Two possible world states $\mathcal{W} = \{L, H\}$ reflect the overall tendency of the preferences, 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 cannot be observed directly 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. Each agent receives a signal in $\mathcal{S} = \{l, h\}$ related to the world state. We use m to denote a generic signal and S_i to denote the random variable representing the signal that agent i receives. We assume the signals are independent and have identical distributions conditioned on the world state. $P_{mk} = \Pr[S_i = m \mid W = k]$ is the probability that an agent receives signal m under

world state k . The signal distributions $((P_{hH}, P_{lH}), (P_{hL}, P_{lL}))$ are common knowledge. The signals are positively correlated to the world states. We assume $P_{hH} > P_{hL}$ and $P_{lH} < P_{lL}$. On the other hand, we allow biased signals and DO NOT assume $P_{hH} > P_{lH}$ or $P_{hL} < P_{lL}$, which is a stronger assumption.




Fig. 1. The information structure of the world state and private signals.

Utility and Types of Agents. Each agent i has a utility function on the true world state and the outcome of the vote. Here we assume agents have 0-1 utility functions to convey the main idea of our work while hiding much of the complexity, Formally, we have $v_i : \mathcal{W} \times \{\mathbf{A}, \mathbf{R}\} \rightarrow \{0, 1\}$. The extension to the setting with general integer utility functions is in Section 6.

Agents can be categorized into three types by their preferences. *Predetermined* agents always prefer the same alternative, and *contingent* agents have preferences depending on the world state. Predetermined agents can be further categorized into *friendly* and *unfriendly* agents based on the alternative they prefer. The utility function of each type of agents is in Table 2.

Type	$v_i(H, \mathbf{A})$	$v_i(L, \mathbf{A})$	$v_i(H, \mathbf{R})$	$v_i(L, \mathbf{R})$
Friendly	1	1	0	0
Unfriendly	0	0	1	1
Contingent	1	0	0	1

Table 2. Utility of agents.

Let F, U, C be the set of three types of agents respectively, and 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*, the alternative favored by the majority given the world state is known to all the agents. We assume that neither α_F nor α_U is larger than 0.5. Otherwise, friendly or unfriendly agents can dominate the voting outcome, and the alternative favored by the majority will be independent of the world state. Therefore, in the rest of this paper, \mathbf{A} is the informed majority decision when the world state is H , and \mathbf{R} is the informed majority decision when the world state is L .

Example 2 (Information Structure). Consider the COVID-19 policy-making scenario. Suppose $n = 20$ voters are deciding whether to accept (\mathbf{A}) or reject (\mathbf{R}) the policy. World state H represents a viral pandemic, and world state L represents the opposite. The common prior on the world state is $P_H = 0.6$, and $P_L = 0.4$. The signals are noisy information about the world state. Suppose the community exaggerates the pandemic, and the signal is biased. The distribution is $P_{hH} = 0.8$, $P_{lH} = 0.2$, $P_{hL} = 0.6$, and $P_{lL} = 0.4$. An agent receives an h signal with probability 0.8 when the world state is H and probability 0.6 when the state is L .

The majority agents has a small advantage on the fraction with $\alpha = 0.6$. That is, there are 30 majority agents and 20 minority agents.

Among 20 agents, there are 5 friendly agents, 6 unfriendly agents, and 9 contingent agents. Neither friendly nor unfriendly agents consist of the majority. Therefore, the informed majority decision is **A** when the world state is H and **R** when the world state is L . \square

Sequence of Environments. An environment \mathcal{I} is a set of parameters of a voting game, which contains the agent number n , the world state prior distribution (P_L, P_H) , the signal distributions (P_{hH}, P_{lH}) and (P_{hL}, P_{lL}) , and the approximated fraction of each type of agents $(\alpha_F, \alpha_U, \alpha_C)$.

In this paper, we study the problem asymptotically. Therefore, we define a sequence of environments in which the number of voters increases while most other parameters remain the same. Let $\{\mathcal{I}_n\}_{n=1}^\infty$ (or $\{\mathcal{I}_n\}$ for short) be a sequence of environments, where each \mathcal{I}_n is an environment of n agents. The environments in a sequence share the same prior distributions, signal distributions, and approximated fractions of each type.

2.1 One round voting

One-round voting. In the one-round voting game, every agent casts a vote for either **A** or **R**. The alternative getting more than a 0.5 fraction of votes becomes the winner.

Strategy in the one-round voting game. We use $\hat{\sigma}$ to denote the strategy of an agent in the one-round voting game. $\hat{\sigma}$ is a mapping from an agent's signal to a probability distribution on $\{\mathbf{A}, \mathbf{R}\}$ representing her action. A strategy is pure if the vote is deterministic. A strategy profile $\hat{\Sigma}$ is the vector of strategies of all agents. A strategy profile is pure if it contains only pure strategies and symmetric if all agents with equal utility function play the same strategy. When agents have 0-1 utilities, symmetry is equivalent to agents of the same type playing the same strategy.

Definition 2.1 (Informative Voting). An agent votes informatively if he/she always votes for **A** when his/her signal is h and always votes for **R** when his/her signal is l .

Expected utility. Let $\lambda_k^{\mathbf{A}}(\hat{\Sigma})$ ($\lambda_k^{\mathbf{R}}(\hat{\Sigma})$, respectively) be the (ex-ante, before agents receiving their signals) probability that **A** (**R**, respectively) becomes the winner in the one round voting when the world state is k . The expected utility of an agent i exclusively depends on $\lambda_k^{\mathbf{A}}(\hat{\Sigma})$ and $\lambda_k^{\mathbf{R}}(\hat{\Sigma})$:

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

Sequence of Strategy Profiles. We define a sequence of strategy profiles $\{\hat{\Sigma}_n\}_{n=1}^\infty$ on an environment sequence $\{\mathcal{I}_n\}$, where for each n , $\hat{\Sigma}_n$ is a strategy profile of a one-round voting game with parameters in \mathcal{I}_n . We do not have additional assumptions about the agents. Therefore, for different n , the strategies of agents can be drastically different.

ε -strong Bayes Nash Equilibrium (BNE). Due to the nature of the majority rule, deviation from a single agent can unlikely change the outcome. Therefore, we take group strategic behavior of agents in to consideration and adopt the solution concept of (approximated) strong Bayes Nash equilibrium. In an ε -strong Bayes Nash Equilibrium, no group of agents can increase their utilities by more than ε through deviation. A strategy profile $\hat{\Sigma} = (\hat{\sigma}_1, \hat{\sigma}_2, \dots, \hat{\sigma}_n)$ is an ε -strong Bayes Nash Equilibrium (ε -strong BNE) if there does not exist a subset of agents D (the *deviating group*) and a strategy profile $\hat{\Sigma}' = (\hat{\sigma}'_1, \hat{\sigma}'_2, \dots, \hat{\sigma}'_n)$ (the *deviating strategy profile*) such that

- (1) $\hat{\sigma}_i = \hat{\sigma}'_i$ for all $i \notin D$;
- (2) $u_i(\hat{\Sigma}') \geq u_i(\hat{\Sigma})$ for all $i \in D$; and

(3) there exists $i \in D$ such that $u_i(\hat{\Sigma}') > u_i(\hat{\Sigma}) + \varepsilon$.

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.

3 Equilibria in One-Round Voting

Han et al. [2023] show that a strategy profile where predetermined agents play their dominant strategies leads to the informed majority decision with high probability if and only if it is an ε -strong BNE. This implies the informed majority decision is (approximately) reached for every equilibrium in the one-round voting game. However, the equilibria in the one-round voting game can be complicated and unnatural. In this section, we show scenarios where natural strategy profiles fail to reach high fidelity (and thus are not approximate strong BNEs according to the if-and-only-if characterization) and how complex the existing equilibria are.

Firstly, in games with only contingent agents (so all agents have the same object), NO pure symmetric equilibria exist in the one round voting game in ANY environment with biased signals.

PROPOSITION 1. *Suppose $\alpha_C = 1$. Then for any sequence of environments $\{\mathcal{I}_n\}$ such that $P_{hL} > 0.5$ or $P_{lH} > 0.5$, no sequence pure symmetric strategy profiles $\{\hat{\Sigma}_n\}$ of the one-round voting game satisfies that $\hat{\Sigma}_n$ is a ε -strong BNE for every n with $\varepsilon \rightarrow 0$.*

PROOF. Without loss of generality, we consider the case that $P_{hL} > 0.5$, and the reasoning for the case that $P_{lH} > 0.5$ is similar. There are in total four pure symmetric strategy profiles: all agents always vote for **A**, always vote for **R**, vote informatively, and vote against their private signals (vote for **A** when receiving signal ℓ and for **R** when receiving signal h). Always voting for an alternative has a probability of P_H or P_L to miss the informed majority decision and cannot be an equilibrium. When $P_{hL} > 0.5$, agents are more likely to receive h than ℓ on world state L . Therefore, by applying the Hoeffding inequality, informative voting does not lead to the informed majority decision in the world state L with high probability and cannot be an equilibrium. Similarly, voting against signals cannot be an equilibrium as $P_{hH} > 0.5$. \square

More characterizations that natural strategies — for example, sincere voting — fail to be equilibria in one-round voting are in Appendix 8.

In the environments where pure symmetric strategies fail to be equilibria, reaching the informed majority decision requires agents to artificially “shift” the probabilities in the strategy. For example, consider a biased setting where the common prior is $P_H = P_L = 0.5$, and the signal distribution is $P_{hH} = 0.9, P_{lH} = 0.1, P_{hL} = 0.7$, and $P_{lL} = 0.3$. An equilibrium under this setting is a symmetric mixed profile where agents with signal l always vote for **R**, and agents with h vote for **A** with probability 0.6. Under this strategy profile, the expected fraction of **A** votes under world H is $0.6 \times 0.9 = 0.54 > 0.5$ and the expected fraction of **A** votes under world L is $0.6 \times 0.7 = 0.42 < 0.5$. Alternatively, instead of playing a mixed strategy, agents can play a pure asymmetric profile where 40% of the agents always vote for **R**.

These “successful” strategy profiles, however, are often too complex for the agents to agree on. For the mixed strategy profile, even if agents are clever enough to aware that a probability shift needs to be done, different agents may choose different probability shifts, which may still lead to the undesirable outcome. For the pure asymmetric profile, without an intrinsic incentive, it is difficult and unclear for identical agents to coordinate on who cast the “opposite vote” **R** when receiving signal h .

4 Two-round Voting and Equilibria

In this section, we introduce our two-round voting mechanism. One-round and two-round voting share the same environment, including the number of agents, distributions, the fraction of each type, and utilities. Their difference lies in the game form and the strategies.

Two-round Voting. Agents play a two-round anonymous voting game where the winner is completely determined by the second round. The five steps are shown in Algorithm 1.

ALGORITHM 1: Two-round voting mechanism

- 1: All the agents receive their signals.
 - 2: **First round** (Polling stage): agents cast the first round vote for **A** or **R**.
 - 3: Agents observe the number of votes for **A** in the first round denoted by $x \in \{0, 1, \dots, n\}$.
 - 4: **Second Round** (Voting stage): agents cast the second round vote for **A** or **R**.
 - 5: The alternative that gets more than a 0.5 fraction of votes in the second round becomes the winner.
-

Strategy. A strategy in a two-round voting game σ_i is a pair (σ_i^1, σ_i^2) . The first-round strategy σ_i^1 is a mapping from agent i 's signal to a distribution on $\{\mathbf{A}, \mathbf{R}\}$ denoting her action in the first round. The second-round strategy σ_i^2 is a mapping from agent i 's signal AND the result of the first-round vote x to a distribution on $\{\mathbf{A}, \mathbf{R}\}$ denoting her action in the second round. An agent i votes *informatively* in the first round if he/she always votes for **A** when his/her signal is h and always votes for **R** when his/her signal is l . A strategy profile Σ is the vector of strategies of all agents. The definition of pure and symmetric strategy profile follows that in the one-round voting.

Sequence of Strategy Profile. We define a sequence of two-round voting strategy profiles $\{\Sigma_n\}_{n=1}^{\infty}$ on an environment sequence $\{\mathcal{I}_n\}$, where for each n , Σ_n is a strategy profile of a two-round voting game with parameters in \mathcal{I}_n .

Expected Utility and Fidelity. 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 **A** (**R**, respectively) becomes the winner when the world state is k .

The expected utility in the two-round voting is in the same form as that in the one-round voting.

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

Fidelity is the likelihood that the informed majority decision is reached. In the two-round voting game, the fidelity of a strategy profile Σ is as follows.

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

For contingent agents with 0-1 utility, their expected utility function $u(\Sigma) = A(\Sigma)$.

An ε -strong Bayes Nash Equilibrium in the two-round voting has the same definition as that in the one-round voting except that the components (strategies, expected utilities) are those in the two-round voting game.

Some Remarks. The game of two-round voting is not exactly a standard *Bayesian game*, as the strategy of an agent depends not only on the signal received but also on the outcome of the first round. Our definitions of strategies, utility functions, and equilibria are straightforward and natural extensions of the standard Bayesian games.

Another natural formulation for the two round-voting is an *extensive-form game*. We believe this is an interesting direction for future work. However, standard equilibrium solution concepts such

as *subgame perfect Nash equilibrium* and *perfect Bayesian equilibrium* fail to capture the feature of group deviations. The first step in this future direction is to propose a reasonable equilibrium concept that extends strong Bayes Nash equilibrium to extensive-form games.

4.1 Equilibria in Two Round Voting

We show that the two-round voting mechanism is indeed a simple mechanism with a natural equilibrium and leads to the informed majority decision. Our first result shows that every ε -strong BNE in the two-round voting mechanisms with ε converges to 0 has fidelity converging to 1. This guarantees that the mechanism can always achieve an informed majority decision under the strategic behavior of agents.

THEOREM 1 (All ε -BNE are good). *Let $\{\Sigma_n\}$ be a sequence of profiles in the two-round voting game. If for any n , Σ_n is an ε -strong BNE with $\varepsilon = o(1)$, then $\lim_{n \rightarrow \infty} A(\Sigma_n) = 1$.*

Moreover, we show the existence of a natural equilibrium in the two-round mechanism. While the classical informative voting or sincere voting may not be an equilibrium in the one-round voting game, their combination becomes an equilibrium as follows.

“Informative + Sincere” voting. All the friendly agents always vote for **A** in both rounds, and all the unfriendly agents always vote for **R** in both rounds. All the contingent agents vote informative in the first round and vote sincerely based on their observation of the first round outcome. That is, conditioned on the first-round strategies and the first-round outcome x , a contingent agent i with private signal m updates his/her belief on the world state. Then the agent votes for **A** if and only if he/she believes the world state is more likely to be H , i.e., $\Pr[H \mid m, x] \geq \frac{1}{2}$.

THEOREM 2 (“Informative + sincere” equilibrium). *For any environment sequence $\{\mathcal{I}_n\}$, the sequence of the “informative + sincere” two-round voting strategy profiles $\{\Sigma_n^\dagger\}$ satisfies that (1) $\lim_{n \rightarrow \infty} A(\Sigma_n^\dagger) = 1$, and (2) for every n , Σ_n^\dagger is a ε -strong BNE where ε converges to 0 as n goes to ∞ .*

The rationale behind the “informative + sincere” strategy is straightforward—by observing the outcome of the first-round informative voting, every contingent agent knows the private information of all contingent agents anonymously, which informs them of the correct world state with high probability. Therefore, in the second round, they are nearly certain of their preferred alternative and vote accordingly.

While agents may not necessarily perform Bayesian updates in the votes, we show that the intuition of “sharing information and getting updated” derives a large class of “informative + threshold” equilibria. Therefore, as agents still follow this “informative + threshold” pattern, the informed majority can be reached.

In an “informative + threshold” strategy profile, all the predetermined agents vote for their preferred alternative in both rounds; every contingent agent i votes informatively in the first round and votes for **A** in the second round if and only if the first round outcome x exceeds some threshold x_i^* . For the “informative + sincere” voting, the threshold x^* is determined by the constraint $\Pr[H \mid m, x] \geq \frac{1}{2}$ as agents apply the Bayes Theorem to update their belief. The explicit form is

$$x_i^* = n_F + \frac{\log \frac{P_L}{P_H} + n_C \cdot \log \frac{P_{tL}}{P_{tH}}}{\log \frac{P_{hH}P_{tL}}{P_{tH}P_{hL}}}.$$

Example 3. We calculate the “informative + sincere” strategy under the setting in Example 2. In this environment, $n_F = 5$, $n_C = 9$, $P_H = 0.6$, $P_{hH} = 0.8$, and $P_{hL} = 0.6$. Therefore, $x_i^* = \frac{\log \frac{2}{3} + 9 \times \log 2}{\log \frac{8}{3}} \approx 10.95$. Therefore, the contingent voters who vote sincerely in the second round will vote for **A** if and only if there are at least 11 votes for **A** in the first round. \square

Additionally, “informative + threshold” class also includes more “natural” strategy profiles, such as the “informative + surprisingly popular” strategy profile.

Example 4. The “informative + surprisingly popular” strategy shares the similar idea with the “surprisingly popular” mechanism [Prelec et al., 2017]. The threshold of every agent i is the prior expectation on the first-round outcome x , which is $x^* = n_F + n_C \cdot (P_H \cdot P_{hH} + P_L \cdot P_{hL})$. This comes from that the n_F friendly agents always vote for A, and the expected share of n_C contingent agents voting for A is P_{hH} under world state H and P_{hL} under world state L . When the number of votes for A exceeds the expectation, A becomes the “surprisingly popular” alternative in the first round and is elected the winner in the second round. Otherwise, R becomes the “surprisingly popular” alternative. \square

The “informative + threshold” strategy profiles share the merit the “informative + sincere” equilibrium has. Every agent votes deterministically. When agents with equal utility functions adopt the same threshold, they are also symmetric. Most importantly, every “informative + threshold” strategy profile in which all thresholds lie between the expected first-round outcome conditioned on world state H and on world state L is an equilibrium that leads to the informed majority decision.

Definition 1. We say an sequence of “informative + threshold” strategies is constantly separable if there exists a $n_0 > 0$ and a constant $\delta > 0$ such that for every $n \geq n_0$ and every agent i , $n_F + n_C \cdot P_{hL} + \delta \cdot n \leq x_i^* \leq n_F + n_C \cdot P_{hH} - \delta \cdot n$.

As illustrated in Figure 2, the expected votes for A in the first round is $n_F + n_C \cdot P_{hH}$ under world state H and $n_F + n_C \cdot P_{hL}$ under world state L .

THEOREM 3 (“Informative + threshold” equilibrium). *Under any environment sequence $\{\mathcal{I}_n\}$, every constantly separable sequence of profiles $\{\Sigma_n^*\}$ satisfies that (1) $\lim_{n \rightarrow \infty} A(\Sigma_n^*) = 1$, and (2) for every n , Σ_n^* is a ε -strong BNE where ε converges to 0 as n goes to ∞ .*




Fig. 2. Illustration of the range of the second round threshold for an “informative + threshold” equilibrium.

Remark on Theorem 1. Theorem 1 states that any equilibrium leads to nearly optimal fidelity. Unlike the one-round voting, the other direction—optimal fidelity implying equilibria—does not hold in the two-round voting game, even when predetermined agents only play their dominant second-round strategies. Intuitively, some supportive behaviors of the predetermined agents in the first round can also constitute a strategy profile that leads to optimal fidelity, while predetermined agents have incentives to deviate from such behaviors to more self-interested strategies.

For example, friendly and unfriendly agents voting informatively in the first round can help contingent agents infer the world state and lead to an informed majority decision. However, this does not form an equilibrium, as friendly agents have incentives to always vote for A in the first

round and convince contingent agents that the world state is H . Similarly, unfriendly agents have incentives to vote for R .

4.2 Proof Sketch for Theorem 1

Suppose $\{\Sigma_n\}$ is a sequence of strategy profiles whose fidelity does not converge to 1. We show that there exists a different sequence of profiles $\{\Sigma'_n\}$ such that agents have incentives to deviate from Σ_n to Σ'_n for infinitely many n .

We define a strategy profile in the two-round voting as *regular* if all friendly agents always vote for A and all unfriendly agents always vote for R in the *second* round. We first give the proof when $\{\Sigma_n\}$ is a regular profile sequence, which characterizes the main idea of the proof while reducing complication, and then extend the proof to non-regular profiles. The full proof can be found in Appendix 8.

4.2.1 Proof for regular profiles. We take an infinite set $\mathcal{N} \subseteq \mathbb{N}_+$ and a constant $\delta > 0$ such that for all $n \in \mathcal{N}$, the fidelity of Σ_n does not exceed $1 - \delta$. Let the deviating group D be all the contingent agents, who prefer A in world state H and R in world state L . Therefore, if the contingent agents switching from Σ_n to Σ'_n increases the fidelity, they get higher expected utility and are incentivized to do so.

In a regular Σ_n , all the predetermined agents vote for their preferred alternative in the second round. If all the contingent agents vote for the same alternative in the second round, this alternative will be the winner. Therefore, our construction of Σ'_n aims to achieve the result that contingent agents reveal the world state with high probability in the first round and vote for the preferred agents in the second round.

Contingent agents can (collectively, with high probability) distinguish the world states because signal distributions are different in different states of the world. For all the contingent agents, the expected number of h signals is $P_{hH} \cdot n_C$ in world state H and $P_{hL} \cdot n_C$ in world state L . If all the contingent agents vote informatively in the first round, then there will be likely more votes for A in world state H than in world state L and more R votes in world state L than in world state H . By the Hoeffding inequality, the number of votes will concentrate to the expectation with high probability as n increases. Therefore, contingent agents can play a “threshold” strategy in the second round: vote for A together if the first round result x exceeds the threshold and vote for R together otherwise. The threshold should be set between $P_{hH} \cdot n_C$ and $P_{hL} \cdot n_C$.

However, this strategy still must overcome one difficulty: strategies of predetermined agents in the first round are not guaranteed to be “regular” and may offset the difference in the votes of contingent agents. For example, predetermined agents can vote for R more often in world state H than in world state L in Σ_n so that x has a smaller expectation in world state H than in world state L . However, contingent agents can “reverse” their votes and observations accordingly: they vote opposite to their signal in the first round and vote for A if x is below a threshold. In this way, they can still distinguish the different world states.

Now we explicitly present the strategy profile Σ'_n . The deviators D are exactly all the contingent agents. Let $\varphi_H(\Sigma'_n)$ and $\varphi_L(\Sigma'_n)$ be the expectation of the first-round outcome x of Σ'_n conditioned on the world state H and L respectively.

- All predetermined agents (they do not deviate) play the same strategy as in Σ_n . This includes that they play their dominant second-round strategies.
- If more votes for A from predetermined agents are expected in world state H than in L in the first round, all contingent agents vote informatively in the first round. For the second round, they vote for A if $x \geq \frac{\varphi_H(\Sigma'_n) + \varphi_L(\Sigma'_n)}{2}$ and vote for R otherwise.

- Otherwise, contingent agents “reverse” their strategy. In the first round, every contingent agent votes for **A** if receiving l and for **R** if receiving h . In the second round, they vote for **A** if $x < \frac{\varphi_H(\Sigma'_n) + \varphi_L(\Sigma'_n)}{2}$ and **R** otherwise.

As shown in the previous analysis, under Σ'_n , contingent agents can distinguish the world state with high probability via the first-round vote, guaranteed by the Hoeffding Inequality. Therefore, the fidelity $A(\Sigma'_n)$ converges to 1 as n goes to infinity. In this way, for all sufficiently large n such that $A(\Sigma_n)$ is not close to 1, by deviating from Σ_n to Σ'_n , the expected utilities of contingent agents increase by a constant. Therefore, contingent agents have incentives to deviate, and Σ_n is NOT an ε -strong BNE for some constant ε .

4.2.2 Proof Sketch for non-regular sequences. We adopt the same idea to construct a strategy profile where the deviating group reveals the world state by the first-round vote and votes for the preferred agent in the second round. Then we show that as the fidelity converges to 1, the group has incentives to do so.

However, for non-regular profiles, the deviating group containing only contingent agents may not work, as there is no assumption on the behavior of predetermined agents in the second round. (For example, if $\alpha_C < 0.5$, and all the predetermined agents always vote for **A**, then the winner will always be **A** regardless of the votes from contingent agents.) Fortunately, the following lemma shows that there will always be a type of predetermined agents have incentives to deviate together to reach the informed majority decision.

LEMMA 1. *Let $\{\Sigma_n\}$ be a sequence of profiles such that there exists a constant δ and an infinite set \mathcal{N} such that for all $n \in \mathcal{N}$, $A(\Sigma_n) \leq 1 - \delta$. Then there exists an infinite set $\mathcal{N}' \subseteq \mathcal{N}$ such that for each $n \in \mathcal{N}'$, at least one of the following holds for all sequences of profiles $\{\Sigma'_n\}$ with $\lim_{n \rightarrow \infty} A(\Sigma'_n) = 1$:*

- (1) *For every friendly agent i_1 , $u_{i_1}(\Sigma'_n) - u_{i_1}(\Sigma_n) \geq 0$.*
- (2) *For every unfriendly agent i_2 , $u_{i_2}(\Sigma'_n) - u_{i_2}(\Sigma_n) \geq 0$.*

Moreover, for a fixed n and Σ_n , either 1 holds for all Σ'_n or 2 holds for all Σ' .

Proof Sketch for Lemma 1. The fidelity $A(\Sigma_n)$ is a weighted sum of the likelihood that **A** wins in world state H (denoted as $\lambda_H^A(\Sigma_n)$) and the likelihood that **R** wins in world state L (denoted as $\lambda_L^R(\Sigma_n)$). From Σ_n to Σ'_n , both likelihoods increase. If $\lambda_H^A(\Sigma_n)$ increases more, the friendly agents get utility increase; and if $\lambda_L^R(\Sigma_n)$ increases more, the unfriendly agents get utility increase. Moreover, as $A(\Sigma'_n)$ is close to 1, which likelihood increases more solely depends on Σ_n . Therefore, which type of agents gets utility increase does not depend on the deviating strategy Σ'_n .

In this way, we consider two deviating strategies. In the first strategy Σ_F , all the friendly and contingent agents vote informatively in the first round and play a threshold strategy in the second round. In the second strategy Σ_U , all the unfriendly agents and all the contingent agents vote informatively in the first round and play a threshold strategy in the second round. Since $\alpha_F < 0.5$ and $\alpha_U < 0.5$, if contingent agents and one type of predetermined agents vote for the same alternative, that alternative will be the winner. Therefore, with similar reasoning to the regular case, both Σ_F and Σ_U have fidelity converging to 1. Applying Lemma 1, either all friendly agents or all unfriendly agents get the expected utility increase via deviation. Moreover, the type of agents with utility gain is the same in Σ_F and Σ_U . Contingent agents will also get a constant increase in expected utility due to the increase in fidelity. Therefore, for each Σ_n whose fidelity is not close to 1, agents have incentives to deviate to either Σ_F or Σ_U . Therefore, Σ_n is NOT an ε -strong BNE with constant ε .

4.3 Proof Sketch for Theorem 2 and Theorem 3.

As the “informative + sincere” strategy profile is also an “informative + threshold” strategy profile, it suffices to display the proof for Theorem 3 and shows that the sincere threshold satisfies the constantly separable constraint.

The proof of Theorem 3 consists of two parts: Lemma 2 shows that the fidelity converges to 1, and Lemma 3 shows that the profiles are $o(1)$ -strong BNE. The full proof is in Appendix 8.

LEMMA 2. *For any sequence of instance and any sequence of “informative + threshold” strategy profiles Σ , if there exists a constant $\delta > 0$ such that for all n and all agent i , $n_F + n_C \cdot P_{hL} + \delta \cdot n \leq x_i^n \leq n_F + n_C \cdot P_{hH} - \delta \cdot n$, then $A(\Sigma)$ converges to 1.*

Lemma 2 follows similar reasoning in Theorem 1 of showing the deviating strategy profile have fidelity converging to 1. Firstly, the threshold x^* is guaranteed to be between the expectation of the first-round vote conditioned on the world state being H and being L . By the Hoeffding inequality, the result of the first round concentrates on the expectation with high probability. Therefore, with probability converging to 1, the first-round outcome is larger than all the thresholds under world state H and lower than all the thresholds under world state L . This informs all the contingent agents of the world state almost surely, and they will vote for the informed majority decision with high probability.

LEMMA 3. Σ_n^* is an ε -strong BNE with $\varepsilon = o(1)$.

The proof of Lemma 3 follows a similar idea to the proof by Schoenebeck and Tao [2021, Theorem 3.3]. For arbitrary n , and a strategy profile Σ_n , let $\varepsilon = 2B(B+2)(1 - A(\Sigma_n))$, where B is the upper bound of the utility. Then, for any deviating strategy profile Σ'_n , we show that the deviation will not succeed. Therefore, Σ_n is an ε -strong BNE. Since $A(\Sigma_n^*)$ converges to 1, ε converges to 0.

LEMMA 4 (INFORMAL). *For any strategy profile Σ , a deviating group D that have incentives to deviate from Σ to another strategy profile Σ' , under the approximation parameter $\varepsilon = 6(1 - A(\Sigma))$ for ε -strong BNE, contains either only friendly agents or only unfriendly agents.*

The proof of Lemma 4 consists of two cases. If $1 - A(\Sigma') < 2 \cdot (1 - A(\Sigma))$, the fidelities of both profiles are close to 1, and the expected utilities of agents will have no significant difference in two profiles. Therefore, no agents can get more than ε expected utility gain by deviation. Otherwise, we show that the deviating group contains only one type of predetermined agents. Firstly, contingent agents, whose expected utilities exactly equals to the fidelity, are excluded from the deviating group D , as their expected utilities decrease with the fidelity. Then, friendly agents and unfriendly agents cannot be in D simultaneously, as an expected utility gain of more than ε on one side implies an expected utility decrease on the other side. The formal statement of Lemma 4 is in Appendix 8.

Finally, we show that D contains only friendly agents or only unfriendly agents cannot succeed. Without loss of generality, suppose D contains only friendly agents. If friendly agents deviate in the first round, the result of the first-round voting x will be less likely to reach the threshold, and contingent agents are more likely to vote for R . If friendly agents deviate in the second round, there will be strictly fewer agents voting for A in expectation, and A will be less likely to be the winner. Therefore, any deviation cannot succeed, and Σ_n^* is a ε -strong BNE with ε converges to 0.

Finally, we show that the “informative+sincere” threshold is between the two expectation.

LEMMA 5. *The sequence of “informative + sincere” strategy profiles is constantly separable.*

Lemma 5 shares the same idea with Lemma 2. Recall that threshold x^* is the boundary for $\Pr[H \mid x, m] \geq \frac{1}{2}$. When x is close to the upper bound (expected utility when the world state is H), $\Pr[H \mid x, m]$ is close to 1; and when x is close to the lower bound (expected utility when the world state is L), $\Pr[H \mid x, m]$ is close to 0. Therefore, the threshold must be in between.

5 Two-Round Voting is More Powerful

Section 3 and Section 4 form a sharp contrast between one-round and two-round voting equilibria. In the one-round voting, any environment with a biased signal structure has no simple and natural strategies form equilibria (and, equivalently, reach the informed majority decision). In the two-round voting, Theorem 2 guarantees a natural strategy profile to reach the informed majority decision and be an equilibrium.

In this section, we show that in any sequence of environments where there exists a one-round equilibrium (where predetermined agents play their dominant strategy), there always exists a two-round equilibrium that will not be more “complicated” than the one-round equilibrium. Precisely, in this two-round equilibrium, agents disregard the first-round voting and play exactly as the one-round equilibrium in the second round. A one-round voting strategy profile is *regular* if all predetermined agents play their dominant strategies.

THEOREM 4. *(Informal) Under the same sequence of environments, for any regular one-round voting strategy profile sequence being $o(1)$ -strong BNE, any two-round voting profile sequence, where agents play exactly this one-round strategy in the second round, is also an $o(1)$ -strong BNE.*

By not exploiting the information revelation in the first round, the two-round voting game was reduced to the one-round voting game. Therefore, we claim that two-round voting is a more powerful tool than one-round voting, as it can reach an informed majority decision whenever one-round voting can reach it, and it provides natural equilibrium in every environment including those where one-round voting fails to have a natural solution.

Proof Sketch. The proof proceeds in three steps. Firstly, the fidelity of the two-round profile sequence converges to 1, because agents in the second round behave exactly as they are playing a one-round game, and the fidelity of the one-round profile converges to 1. Secondly, by applying Lemma 4, a deviating group either contains only friendly agents or only unfriendly agents. Finally, such a deviating group cannot succeed, similar to Lemma 3. A more technical description of Theorem 4 and the full proof is in Appendix 8.

6 General Valuation Utility.

In this section, we extend our theoretical results to a setting with more general utility functions.

General Valuation utility. The utility function of each agent $v_i : \mathcal{W} \times \{\mathbf{A}, \mathbf{R}\} \rightarrow \{0, 1, 2, \dots, B\}$, where $B > 0$ is the integer upper bound. In our model, \mathbf{A} is more welcomed in H , and \mathbf{R} is more welcomed in L : $v_i(H, \mathbf{A}) \geq v_i(L, \mathbf{A})$, and $v_i(L, \mathbf{R}) \geq v_i(H, \mathbf{R})$. For an agent i , if i is a friendly agent, $v_i(H, \mathbf{A}) \geq v_i(L, \mathbf{A}) > v_i(L, \mathbf{R}) \geq v_i(H, \mathbf{R})$; if i is an unfriendly agent, $v_i(L, \mathbf{R}) \geq v_i(H, \mathbf{R}) > v_i(H, \mathbf{A}) \geq v_i(L, \mathbf{A})$; and if i is a contingent agent, $v_i(H, \mathbf{A}) > v_i(H, \mathbf{R})$ and $v_i(L, \mathbf{R}) > v_i(L, \mathbf{A})$.

Under the general valuation utility setting, agents of the same type have the same ordinal preference, yet their utility functions may be different. For example, the utility function of a contingent agent may not equals to (or be proportional to) the fidelity anymore. All other definitions in the setting remain the same. Specifically, in a symmetric strategy profile, agents with equal utility function play the same strategy. This allows agents with the same ordinal preferences to play differently.

As the utility diverges, a sincere voter does not always vote for the alternative that is more likely to be the informed majority decision. Instead, they vote for the alternative that brings them a higher expected utility.

Definition 2 (Sincere Voting in the second round). A sincere agent votes as if he/she is making an individual decision. The agent compares his/her expected utility conditioned on his/her private

signal m , the first-round outcome x , and the event that A/R becomes the winner, respectively.

$$u_i(\mathbf{A} \mid x, m) = \Pr[W = L \mid x, m] \cdot v_i(L, \mathbf{A}) + \Pr[W = H \mid x, m] \cdot v_i(H, \mathbf{A}).$$

$$u_i(\mathbf{R} \mid x, m) = \Pr[W = L \mid x, m] \cdot v_i(L, \mathbf{R}) + \Pr[W = H \mid x, m] \cdot v_i(H, \mathbf{R}).$$

If $u_i(\mathbf{A} \mid x, m) \geq u_i(\mathbf{R} \mid x, m)$, the agent vote for A; otherwise, the agent vote for R.

Now we extend our results to the general valuation setting. The sharp contrast between the one-round voting and the two-round voting still holds. For one-round voting, instances with NO pure symmetric equilibrium still exist.

PROPOSITION 1’. *There exists a sequence of environments $\{\mathcal{I}_n\}$ such that no sequence pure symmetric strategy profiles $\{\hat{\Sigma}_n\}$ of the one-round voting game satisfies that $\hat{\Sigma}_n$ is a ε -strong BNE for every n with $\varepsilon \rightarrow 0$.*

The environments specified in Proposition 1, where all agents are contingent (and have equal utility functions) and the signal is biased, serves as the proof for Proposition 1’.

In the sharp contrast, all our positive results in the two-round voting holds in the general valuation utility setting.

THEOREM 1’ (All ε -BNE are good). *Let $\{\Sigma_n\}$ be a sequence of profiles in the two-round voting game. If for any n , Σ_n is an ε -strong BNE with $\varepsilon = o(1)$, then $\lim_{n \rightarrow \infty} A(\Sigma_n) = 1$.*

THEOREM 2’ (“Informative + sincere” equilibrium). *For any environment sequence $\{\mathcal{I}_n\}$, the sequence of the “information + sincere” two-round voting strategy profiles $\{\Sigma_n^\dagger\}$ satisfies that (1) $\lim_{n \rightarrow \infty} A(\Sigma_n^\dagger) = 1$, and (2) for every n , Σ_n^\dagger is a ε -strong BNE where ε converges to 0 as n goes to ∞ .*

THEOREM 3’ (“Informative + threshold” equilibrium). *Under any environment sequence $\{\mathcal{I}_n\}$, every constantly separable sequence of profiles $\{\Sigma_n^*\}$ satisfies that (1) $\lim_{n \rightarrow \infty} A(\Sigma_n^*) = 1$, and (2) for every n , Σ_n^* is a ε -strong BNE where ε converges to 0 as n goes to ∞ .*

THEOREM 4’. *(Informal) Under the same sequence of environments, for any regular one-round voting strategy profile sequence being $o(1)$ -strong BNE, any two-round voting profile sequence, where agents play exactly this one-round strategy in the second round, is also an $o(1)$ -strong BNE.*

Remark for Theorem 2’. Under the general valuation setting, as the behavior of sincere agents diverges with the utility, the “sincere” threshold depends on the utility of the agents. Formally, for a contingent agent i ,

$$x_i^* = n_F + \frac{\log \frac{P_L \cdot (v_i(L, \mathbf{R}) - v_i(L, \mathbf{A}))}{P_H \cdot (v_i(H, \mathbf{A}) - v_i(H, \mathbf{R}))} + n_C \cdot \log \frac{P_{FL}}{P_{FH}}}{\log \frac{P_{HH} P_{FL}}{P_{FH} P_{HL}}}.$$

Agents with different utility functions may have different threshold x_i^* , yet this gap remains constant as n grows. Therefore, in large elections, sincere voters adopt different yet close thresholds.

The proofs of Theorem 1’ to 4’ follow a similar reasoning to their 0-1 utility version, yet the general valuation utilities bring new technical challenges. As the preferences vary, the utility change of the agents are less aligned with the change in the fidelity. For example, in Lemma 1, we show that a type of predetermined agents are happy with fidelity increases. However, in the general valuation setting, the Lemma further relies on the assumptions that agents have different sensitivities of utilities under different world state. The friendly agents are more sensitive in world state H ($v_i(H, \mathbf{A}) - v_i(H, \mathbf{R}) > v_i(L, \mathbf{A}) - v_i(L, \mathbf{R})$), while unfriendly agents are more sensitive in world state L . Therefore, they usually have higher utility increase when the fidelity increases compared to the utility decrease when the fidelity decreases, which improve their willingness to deviate to increase expected utility. All the proofs are in Appendix 8

7 Experiments

In addition to our theoretical results, we conducted simulated voting experiments with generative AI agents to compare one-round and two-round voting. Specifically, we inject generative AI bots with different preferences and information drawn from given distributions, then run both one-round and two-round voting under various environment parameters. Our findings show that the generative AI agents achieved the informed majority decision more often under the two-round mechanism than under the one-round mechanism.

7.1 Experimental Setting

Our experiments employed two generative AI models—OpenAI GPT-4o and Deepseek R1, a newly released reasoning model. In each voting scenario, all voters were drawn from the same model.

For each vote, we set $n = 40$ agents, including 5 friendly agents, 3 unfriendly agents, and 32 contingent agents. The prior is fixed at $P_H = 0.5$. We run experiments on two different signal distributions, the unbiased distribution $P_{hH} = 0.8, P_{hL} = 0.2$ and the biased distribution $P_{hH} = 0.8, P_{hL} = 0.6$. We also test a wider range of signal distributions with fewer samples and select the distributions that capture different scenarios anticipated by our theoretical findings on the performance of one-round versus two-round voting. Under the unbiased distribution, both mechanisms admit natural equilibria that achieve the informed majority decision, so we expect strong performance from both. In contrast, the biased distribution, as shown by Proposition 1, admits no pure-symmetric equilibria for one-round voting, whereas “informative + sincere” remains viable in the two-round mechanism—implying better performance in that setting.

We run 50 votes for each set of parameters. In each vote, we first sampled the world state and signals for every agent from the designated distribution. Using these same samples, we then ran both a one-round and a two-round voting procedure. Specifically, we copy n generative AI agents for the one-round voting and another n agents for the two-round voting. In one-round voting, each agent was asked only once for its vote. In two-round voting, each agent voted twice: we first collected every agent’s first-round vote simultaneously, then shared the first-round outcome with all agents (including their own first-round response as a history message) when asking for the second-round vote.

The prompt for AI agents contains the following components.

Voting scenario. We chose a company hiring setting: each AI agent plays the role of a human resource specialist and votes on whether to hire a candidate.

World state and signals. The world state is whether the candidate is qualified, which is not revealed to the agents. Each AI agent is told to have a Good or Bad impression of the candidate, representing its private signal.

Prior and signal distribution. AI agents are informed of the prior and signal distribution in two different ways. In the *accurate distribution* setting, every agent is told the exact distribution (P_H, P_L) and ($P_{hH}, P_{hL}, P_{\ell H}, P_{\ell L}$). In the *vague distribution* setting, the information is transferred into a vague description via a Deepseek R1 chatbot. For example, the distribution $P_H = 0.5, P_{hH} = 0.8$, and $P_{hL} = 0.6$ is converted the following description.

Example 5 (Vague information in the prompt.) Before meeting with the candidate, you thought she had a balanced chance to be qualified. If the candidate is qualified, you will likely have a good impression after meeting with her with a high probability well above half. If the candidate is not qualified, you will be less likely, though still moderately probable, to have a good impression at a likelihood distinctly lower than the qualified scenario but remaining above the halfway mark. \square

Types and preferences. Every agent is told their preferences based on their type.

Voting Mechanisms. Every agent is informed that it will participate in either a one-round or a two-round voting procedure. For the two-round voting, we clarify that the first round does not determine the winner but is publicly observable.

Vote. In both the one-round voting and each round of the two-round voting, every agent is asked to determine and report its best vote (YES or NO), along with a brief explanation of its reasoning.

7.2 Experimental Results

Model	Distribution	Unbiased (0.8, 0.2)		Biased (0.8, 0.6)	
	Information	Accurate	Vague	Accurate	Vague
Deepseek R1	One-round	1	1	0.56	0.52
	Two-round	1	1	0.88	0.78
GPT-4o	One-round	1	1	0.5	0.52
	Two-round	1	1	0.48	0.5

Table 3. The likelihood of the informed majority decision is reached.

In Table 3, we compare how often one-round versus two-round voting achieves the informed majority decision under different signal distributions, information conditions, and generative AI models. When signals are unbiased, both voting mechanisms always reach the informed majority decision. However, their performances diverge significantly under biased signals. In particular, Deepseek R1 agents show a marked accuracy gap: one-round voting frequently misses the informed majority decision when the world state is unqualified (even though agents still receive a "Good" signal with probability 0.6), whereas two-round voting retains good performance under both world states. On the other hand, for the non-reasoning model GPT-4o, there is no significant different between the performance of one-round voting and two-round voting.

In addition to the analysis of the informed majority decision, we also examine the behaviors of the AI agents from their reasoning and analysis and have some interesting observations.

Bayesian Updates with accurate and vague information. Almost every agent performs a Bayesian update to infer the likelihood of each world state based on the signal distribution and the private information it obtains. Even when vague information is given, many agents tend to guess a concrete value for each description and perform a Bayesian update based on the guessed values.

Vote switch in the second round. We observe a universal switch in the second-round votes where a voter casts a different vote to its first-round vote. Examining the analysis of these voters, they usually find considerable votes in the first round that contrast their beliefs, which drives them to switch. On the other hand, there are also agents who believe more in their beliefs and stick to the same vote. An example of the prompts, answers, and reasoning of agents is in Appendix 8.

7.3 Take-away Messages

Our experiments illustrate the scenario where two-round voting is more likely to lead to an informed majority decision than one-round voting played by generative AI agents. These findings can guide future applications of generative AI in voting systems, such as AI-augmented or AI-proxy voting. As AI reasoning capabilities continue to improve, two-round voting will likely become even more effective in reaching informed outcomes. Moreover, our experiments serve as a practical reference for designing real-world experiments that compare different voting rules and explore strategic

behavior. Because generative AI simulations allow for direct preference injection, they offer a valuable tool for preliminary experimentation and design.

8 Conclusion and Future Work

We study the voting problem with two alternatives where voters' preferences depend on a not-directly-observable state variable and propose the two-round voting mechanism where the first round serves as a polling stage and the winning alternative only depends on the outcome of the second round. Our analysis demonstrates that this mechanism can achieve informed majority decisions under strategic behavior through a natural equilibrium combining informative first-round and sincere second-round voting. This sharply contrasts one-round voting mechanisms that often fail to provide both simplicity and effective equilibria simultaneously. Furthermore, we prove that the two-round mechanism supports second-round consistent equilibria, where voter behavior mirrors that of one-round voting scenarios. These properties establish two-round voting as an elegant solution that combines simplicity with robust information aggregation capabilities, ultimately achieving informed majority decisions effectively.

A future direction is to formulate the two-round voting game as an extensive-form game. To incorporate the analysis with group manipulations, we need to first find a reasonable equilibrium concept that extends the strong Bayes Nash equilibrium to the extensive-form game. It would be interesting to see if the corresponding equilibria lead to good election outcomes and compare the equilibria to the ones in this paper.

In this paper, we assume that agents share a common prior on the world state and private signals. If agents have no information about the prior, then, it has been shown in various settings that strategic agents may not be able to make an informed majority decision [Feddersen and Pesendorfer, 1997, Schoenebeck and Tao, 2021]. However, it is still an interesting question of whether we could relax the common knowledge assumption. For example, what if the agents had common knowledge about the distribution of possible world/signal models? In this case, it remains unknown whether one-round or two-round voting mechanisms can still achieve the informed majority decision. We are interested in whether the polling stages in the two-round or multi-round voting can help agents share information and form a common belief so that the correct decision can be made eventually in the last round.

Another possible direction is to conduct human-subject experiments comparing one-round and two-round voting mechanisms, as a further step of our generative AI experiments. We would like to know both if the polling stage indeed helps real-world voters to aggregate their information and preferences more easily, and how likely and when one-round and two-round mechanisms have different outcomes.

References

- David Austen-Smith. 1993. Interested experts and policy advice: Multiple referrals under open rule. *Games and Economic Behavior* 5, 1 (1993), 3–43.
- 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.
- Eyal Baharad, Jacob Goldberger, Moshe Koppel, and Shmuel Nitzan. 2012. Beyond Condorcet: Optimal aggregation rules using voting records. *Theory and decision* 72, 1 (2012), 113–130.
- Marco Battaglini. 2002. Multiple referrals and multidimensional cheap talk. *Econometrica* 70, 4 (2002), 1379–1401.
- Marco Battaglini. 2004. Policy advice with imperfectly informed experts. *The BE Journal of Theoretical Economics* 4, 1 (2004), 0000102202153459631100.
- Marco Battaglini, Rebecca B. Morton, and Thomas R. Palfrey. 2010. The Swing Voter's Curse in the Laboratory. *The Review of Economic Studies* 77, 1 (01 2010), 61–89.
- 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.

- Duncan Black. 1948. On the rationale of group decision-making. *Journal of Political Economy* 56, 1 (1948), 23–34.
- Samir Chopra, Eric Pacuit, and Rohit Parikh. 2004. Knowledge-Theoretic Properties of Strategic Voting. In *Logics in Artificial Intelligence*, José Júlio Alferes and João Leite (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 18–30.
- 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.
- Peter J. Coughlan. 2000. In Defense of Unanimous Jury Verdicts: Mistrials, Communication, and Strategic Voting. *The American Political Science Review* 94, 2 (2000), 375–393.
- Vincent P Crawford and Joel Sobel. 1982. Strategic information transmission. *Econometrica: Journal of the Econometric Society* (1982), 1431–1451.
- Eddie Dekel and Michele Piccione. 2000. Sequential Voting Procedures in Symmetric Binary Elections. *Journal of Political Economy* 108 (2000), 34–55. Issue 1.
- Ignacio Esponda and Emanuel Vespa. 2014. Hypothetical Thinking and Information Extraction in the Laboratory. *American Economic Journal: Microeconomics* 6, 4 (November 2014), 180–202.
- 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.
- Dino Gerardi and Leeat Yariv. 2007. Deliberative voting. *Journal of Economic Theory* 134, 1 (2007), 317–338.
- Jacob K. Goeree and Leeat Yariv. 2011. An Experimental Study of Collective Deliberation. *Econometrica* 79, 3 (2011), 893–921.
- 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.
- 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* (London, United Kingdom). Association for Computing Machinery, New York, NY, USA, 885–905.
- Serguei Kaniovski. 2010. Aggregation of correlated votes and Condorcet's Jury Theorem. *Theory and Decision* 69, 3 (2010), 453–468.
- John Morgan and Phillip C. Stocken. 2008. Information Aggregation in Polls. *American Economic Review* 98, 3 (June 2008), 864–96. <https://doi.org/10.1257/aer.98.3.864>
- 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.
- 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.
- 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.
- H. Peyton Young. 1988. Condorcet's Theory of Voting. *American Political Science Review* 82 (1988), 1231–1244.

A More equilibrium analysis in one-round voting

In this section, we give more examples where natural strategies fail to be equilibria in the one-round voting, including sincere voting (where an agent chooses the alternative maximizing his/her expected utility when this alternative becomes the winner) and the Bayesian strategy (an agent votes for an alternative with the probability of the Bayesian posterior probability that this alternative is the informed majority decision).

We follow the example in Section 3 where all agents are contingent, and the signal distribution is $P_{hH} = 0.9$, $P_{lH} = 0.1$, $P_{hL} = 0.7$, and $P_{lL} = 0.3$. In Section 3 we have shown that any pure symmetric strategy fails to be an equilibrium in this setting. Now we show that sincere voting fails, and the Bayesian strategy is an equilibrium only when common prior is $P_H = P_L = 0.5$,

One natural strategy is to choose the alternative that maximizes the agent's expected utility, assuming this agent is the only voter who decides the outcome. This is called *the sincere strategy*. Suppose the common prior is $P_H = P_L = 0.5$. An agent's posterior belief for world H upon receiving signal h is

$$\Pr(W = H \mid S_i = h) = \frac{\Pr(W = H) \cdot \Pr(S_i = h \mid W = H)}{\Pr(S_i = h)} = \frac{0.5 \times 0.9}{0.5 \times 0.9 + 0.5 \times 0.7} = \frac{9}{16},$$

and the posterior belief for world H upon receiving signal l is

$$\Pr(W = H \mid S_i = l) = \frac{\Pr(W = H) \cdot \Pr(S_i = l \mid W = H)}{\Pr(S_i = l)} = \frac{0.5 \times 0.1}{0.5 \times 0.1 + 0.5 \times 0.3} = \frac{1}{4}.$$

If an agent i receives h , the expected utility of voting for **A** is

$$\Pr(W = H \mid S_i = h) \cdot v_i(H, \mathbf{A}) + \Pr(W = L \mid S_i = h) \cdot v_i(L, \mathbf{A}) = \frac{9}{16}v_i(H, \mathbf{A}) + \frac{7}{16}v_i(L, \mathbf{A}),$$

and the expected utility of voting for **R** is

$$\Pr(W = H \mid S_i = h) \cdot v_i(H, \mathbf{R}) + \Pr(W = L \mid S_i = h) \cdot v_i(L, \mathbf{R}) = \frac{9}{16}v_i(H, \mathbf{R}) + \frac{7}{16}v_i(L, \mathbf{R}).$$

When voting sincerely, the agent receiving signal h will vote for **A** if and only if the expected utility for **A** is higher:

$$\frac{9}{16}v_i(H, \mathbf{A}) + \frac{7}{16}v_i(L, \mathbf{A}) > \frac{9}{16}v_i(H, \mathbf{R}) + \frac{7}{16}v_i(L, \mathbf{R}).$$

Similarly, the agent receiving signal l will vote for **A** if and only if

$$\frac{1}{4}v_i(H, \mathbf{A}) + \frac{3}{4}v_i(L, \mathbf{A}) > \frac{1}{4}v_i(H, \mathbf{R}) + \frac{3}{4}v_i(L, \mathbf{R}).$$

The utility for a contingent agent satisfies that $v_i(H, \mathbf{A}) > v_i(H, \mathbf{R})$ and $v_i(L, \mathbf{R}) > v_i(L, \mathbf{A})$. However, when agents care more about the outcome under world H and are relatively indifferent about the outcome under world L , they will always vote for **A** and it is not favored by them under world L . In our example, if $v_i(H, \mathbf{A}) - v_i(H, \mathbf{R}) > 3 \cdot (v_i(L, \mathbf{R}) - v_i(L, \mathbf{A}))$, agents will vote for **A** even when signal l is received. Thus, the sincere strategy profile does not guarantee the informed majority decision, and it is not a strong BNE by the if-and-only-if characterization.

Perhaps one more natural strategy is to vote according to the posterior, i.e., adopt the mixed strategy with probabilities matching the posterior belief: an agent votes for **A** with probability $\beta_l = \Pr(W = H \mid S_i = l)$ when receiving signal l and probability $\beta_h = \Pr(W = H \mid S_i = h)$ when receiving signal h . In our specific environment, with the prior belief $P_H = P_L = 0.5$, our previous calculations reveal that the strategy is given by $(\beta_l, \beta_h) = (\frac{1}{4}, \frac{9}{16})$. When the actual world is H , the "correct" alternative **A** gets a share of $0.1 \cdot \frac{1}{4} + 0.9 \cdot \frac{9}{16} = 53.125\%$ in expectation; when the actual world is L , the "incorrect" alternative **A** wins a share of $0.3 \cdot \frac{1}{4} + 0.7 \cdot \frac{9}{16} = 46.875\%$ in expectation.

This looks good: the “correct” alternative always receives more than half of the votes in expectation. Some calculations reveal that the informed majority decision can be guaranteed (more formally, the fidelity goes to 1 when $n \rightarrow \infty$) by this kind of matching-posterior strategy profile if the prior belief is half-half (i.e., $P_H = P_L = 0.5$). However, if the prior belief is significantly biased, say, P_H is close to 1, some calculations reveal that the expected fraction of A votes is close to 1 under both world states, in which case the informed majority decision is not secured under the world state L.

B Full Proofs

We abuse the notation for deterministic first/second round strategy such that $\sigma_i^1(m) = \mathbf{A}$ means that i always votes for A (with probability 1) in the first round when her signal is m . Other deterministic votes follow a similar notation.

The proofs are written under the general valuation setting and can be transferred to those under 0-1 utility setting by assigning the utilities.

B.1 Proof of Lemma 4

LEMMA 4 (FORMAL). *Let Σ^* be an arbitrary profile, and let $\varepsilon = 2B(B+2)(1 - A(\Sigma^*))$. If there exists a group of agents D and another strategy profile Σ' such that agents in D have incentives to deviate to Σ' so that (1) $\forall i \in D, u_i(\Sigma') \geq u_i(\Sigma^*)$, and (2) there exists $i \in D$ such that $u_i(\Sigma') > u_i(\Sigma) + \varepsilon$, then either $D \subseteq F$ or $D \subseteq U$.*

In this proof, we use Σ^* to denote Σ_n^* for simplicity. We show that for any other strategy profile Σ' , there does not exist a group of agents wishing to deviate. Let $e = 1 - A(\Sigma^*)$. We know that e converges to 0 as n goes to 1. We take $\varepsilon = 2B(B+2)e$.

PROOF OF LEMMA 4. We consider two different cases of the fidelity of Σ' .

Case 1: $(1 - A(\Sigma')) < (2B+2)e$. In this case, since the fidelity of Σ' is also close to one, two profiles make no significant difference, and no agent can gain more than $\varepsilon = 2B(B+2)e$.

Consider the difference of expected utility in two profiles.

$$\begin{aligned} u_i(\Sigma') - u_i(\Sigma^*) &= P_H \cdot (\lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*)) (v_i(H, \mathbf{A}) - v_i(H, \mathbf{R})) \\ &\quad + P_L \cdot (\lambda_L^R(\Sigma') - \lambda_L^R(\Sigma^*)) (v_i(L, \mathbf{R}) - v_i(L, \mathbf{A})). \end{aligned}$$

Since we know that $A(\Sigma^*) = 1 - e$, we have $\lambda_H^A(\Sigma^*) \geq 1 - \frac{e}{P_H}$ and $\lambda_L^R(\Sigma^*) \geq 1 - \frac{e}{P_L}$. Similarly, for Σ' , we have $\lambda_H^A(\Sigma') \geq 1 - \frac{(2B+2)e}{P_H}$ and $\lambda_L^R(\Sigma') \geq 1 - \frac{(2B+2)e}{P_L}$. Then we can bound the difference between these probabilities.

$$\begin{aligned} \lambda_L^R(\Sigma^*) - \lambda_L^R(\Sigma') &\leq \frac{(2B+2)e}{P_L}, \\ \lambda_L^R(\Sigma') - \lambda_L^R(\Sigma^*) &\leq \frac{e}{P_L}, \\ \lambda_H^A(\Sigma^*) - \lambda_H^A(\Sigma') &\leq \frac{(2B+2)e}{P_H}, \\ \lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*) &\leq \frac{e}{P_H}. \end{aligned}$$

Then we can bound the expected utilities.

- For the friendly agents, $v_i(H, \mathbf{A}) > v_i(H, \mathbf{R})$ and $v_i(L, \mathbf{A}) > v_i(L, \mathbf{R})$. Then

$$\begin{aligned} u_i(\Sigma') - u_i(\Sigma^*) &\leq P_L \cdot \frac{(2B+2)e}{P_L} \cdot B + P_H \cdot \frac{e}{P_H} \cdot B \\ &= 2B(B+2)e. \end{aligned}$$

- For contingent agents, $v_i(H, \mathbf{A}) > v_i(H, \mathbf{R})$ and $v_i(L, \mathbf{A}) < v_i(L, \mathbf{R})$. Then

$$\begin{aligned} u_i(\Sigma') - u_i(\Sigma^*) &\leq P_L \cdot \frac{e}{P_L} \cdot B + P_H \cdot \frac{e}{P_H} \cdot B \\ &= 2Be. \end{aligned}$$

- For the unfriendly agents, we have $v_i(H, \mathbf{A}) < v_i(H, \mathbf{R})$ and $v_i(L, \mathbf{A}) < v_i(L, \mathbf{R})$. Then

$$\begin{aligned} u_i(\Sigma') - u_i(\Sigma^*) &\leq P_L \cdot \frac{e}{P_L} \cdot B + P_H \cdot \frac{(2B+2)e}{P_H} \cdot B \\ &= 2B(B+2)e. \end{aligned}$$

Therefore, no agents can gain more than $\varepsilon = 2B(B+2)e$.

Case 2: $(1 - A(\Sigma')) \geq (2B+2)e$. In this case, we proceed with our proof in two steps.

- (1) Firstly, contingent agents will not deviate, because their expected utility will strictly decrease.
- (2) Secondly, a deviating group cannot contain both friendly and unfriendly agents.

CLAIM 1. *If $(1 - A(\Sigma')) \geq (2B+2)e$, then for any contingent agent i , $u_i(\Sigma') - u_i(\Sigma^*) < 0$.*

Still, recall the difference between the two expected utilities.

$$\begin{aligned} u_i(\Sigma') - u_i(\Sigma^*) &= P_H \cdot (\lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*)) (v_i(H, \mathbf{A}) - v_i(H, \mathbf{R})) \\ &\quad + P_L \cdot (\lambda_L^R(\Sigma') - \lambda_L^R(\Sigma^*)) (v_i(L, \mathbf{R}) - v_i(L, \mathbf{A})). \end{aligned}$$

From $A(\Sigma^*) = e$ and $A(\Sigma') \leq 1 - (2B+2)e$, we have

$$P_L \cdot \lambda_L^R(\Sigma^*) + P_H \cdot \lambda_H^A(\Sigma^*) - P_L \cdot \lambda_L^R(\Sigma') + P_H \cdot \lambda_H^A(\Sigma') \geq (2B+1)e.$$

Then we consider different cases between the λ .

- If $\lambda_H^A(\Sigma') \leq \lambda_H^A(\Sigma^*)$ and $\lambda_L^R(\Sigma') \leq \lambda_L^R(\Sigma^*)$, at least one of them will be strict. Then $u_i(\Sigma') - u_i(\Sigma^*) < 0$.
- If $\lambda_H^A(\Sigma') > \lambda_H^A(\Sigma^*)$, then we have $\lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*) < \frac{e}{P_H}$. On the other hand, $\lambda_L^R(\Sigma') - \lambda_L^R(\Sigma^*) \leq -\frac{(2B+1)e}{P_L}$. Then,

$$u_i(\Sigma') - u_i(\Sigma^*) \leq P_H \cdot \frac{e}{P_H} \cdot B + P_L \cdot \left(-\frac{(2B+1)e}{P_L} \right) \cdot 1 < 0.$$

- If $\lambda_L^R(\Sigma') > \lambda_L^R(\Sigma^*)$, with similar reasoning, we can show that $u_i(\Sigma') - u_i(\Sigma^*) < 0$.

CLAIM 2. *Let i_1 be any friendly agent and i_2 be any unfriendly agent. Then*

- (1) *If $u_{i_1}(\Sigma') - u_{i_1}(\Sigma^*) > \varepsilon$, then $u_{i_2}(\Sigma') - u_{i_2}(\Sigma^*) < 0$.*
- (2) *If $u_{i_2}(\Sigma') - u_{i_2}(\Sigma^*) > \varepsilon$, then $u_{i_1}(\Sigma') - u_{i_1}(\Sigma^*) < 0$.*

We only prove 1, and 2 will follow the same reasoning. Without loss of generality, let $i_1 = 1$ and $i_2 = 2$.

Still, recall the difference between the expected utility.

$$\begin{aligned} u_i(\Sigma') - u_i(\Sigma^*) &= P_H \cdot (\lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*)) (v_i(H, \mathbf{A}) - v_i(H, \mathbf{R})) \\ &\quad + P_L \cdot (\lambda_L^R(\Sigma') - \lambda_L^R(\Sigma^*)) (v_i(L, \mathbf{R}) - v_i(L, \mathbf{A})). \end{aligned}$$

We consider two cases.

Case 1: $\lambda_H^A(\Sigma') \geq \lambda_H^A(\Sigma^*)$. In this case, since $\lambda_H^A(\Sigma^*)$ is already close to 1, the gain in agent 1's expected utility on H 's side will not exceed $P_H \cdot \frac{e}{P_H} \cdot B = eB < \varepsilon$. Therefore, $\lambda_L^R(\Sigma') < \lambda_L^R(\Sigma^*)$. In this case, the expected utility of agent 2 will strictly decrease.

Case 2: $\lambda_H^A(\Sigma') < \lambda_H^A(\Sigma^*)$. In this case, the P_L term is positive for agent 1, while the P_H term is positive for agent 2. For agent 1, we have $v_1(H, \mathbf{A}) \geq v_1(L, \mathbf{A}) > v_1(L, \mathbf{R}) \geq v_1(H, \mathbf{R})$. Then since $u_1(\Sigma') - u_1(\Sigma^*)$ is strictly positive, we must have

$$P_H \cdot (\lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*)) > P_L \cdot (\lambda_L^R(\Sigma') - \lambda_L^R(\Sigma^*)).$$

For agent 2, on the other hand, we have $v_1(L, \mathbf{R}) \geq v_1(H, \mathbf{R}) > v_1(H, \mathbf{A}) \geq v_1(L, \mathbf{A})$. Then there must have $u_2(\Sigma') - u_2(\Sigma^*) < 0$. \square

B.2 Proof of Theorem 3

The proof of Theorem 3 resembles that of Theorem 2. Lemma 2 guarantees the the fidelity. For the equilibrium parts, applying Lemma 4, it suffice to consider a deviating group with only friendly agents or only unfriendly agents. Finally, Claim 3 shows that such deviation cannot succeed.

LEMMA 2. *Under any environment sequence $\{\mathcal{I}_n\}$, every constantly separable sequence of profiles $\{\Sigma^*\}$ satisfies that (1) $\lim_{n \rightarrow \infty} A(\Sigma_n^\dagger) = 1$, and (2) for every n , Σ_n^* is a ε -strong BNE where ε converges to 0 as n goes to ∞ .*

PROOF. The high-level idea of Σ^* is that contingent agents determine the world state with high probability via the first-round voting and vote for the correct alternative in the second round. Since predetermined agents behave regularly, and both α_F and α_U do not exceed 0.5, the alternative that all contingent agents vote for will be the winner. Now we show that both $\lim_{n \rightarrow \infty} \lambda_H^A(\Sigma^*) = 1$ and $\lim_{n \rightarrow \infty} \lambda_L^R(\Sigma^*) = 1$.

For now we fix an arbitrary n . For an agent i , Let X_i be the random variable indicating i 's vote in the first round. $X_i = 0$ means i votes for \mathbf{R} , and $X_i = 1$ means i votes for \mathbf{A} . Then the result of the first round voting x can be represented as $x = \sum_i X_i$. If the first round outcome exceeds $x \geq n_F + n_C \cdot P_{hH} - \delta$, then all the agents i will vote for \mathbf{A} , and \mathbf{A} becomes the winner. Therefore, the probability such that \mathbf{A} becomes the winner conditioned on the world state H satisfies

$$\lambda_H^A(\Sigma^*) \geq \Pr \left[\sum_i X_i \geq n_F + n_C \cdot P_{hH} - \delta \cdot n \mid H \right].$$

Since contingent agent votes informatively in the first round, the expectation of of the first round outcome $E[\sum_i X_i \mid H] = n_F + n_C \cdot P_{hH}$. Then for all $n \geq n_0$, we can apply the Hoeffding inequality to the probability:

$$\begin{aligned} \lambda_H^A(\Sigma^*) &\geq \Pr \left[\sum_i X_i \geq n_F + n_C \cdot P_{hH} - \delta \cdot n \mid H \right] \\ &= \Pr \left[\sum_i X_i - E[\sum_i X_i \mid H] \geq n_F + n_C \cdot P_{hH} - \delta \cdot n - E[\sum_i X_i \mid H] \mid H \right] \\ &= \Pr \left[\sum_i X_i - E[\sum_i X_i \mid H] \geq -\delta \cdot n \mid H \right] \\ &\geq 1 - \exp(2\delta^2 \cdot n). \end{aligned}$$

Similarly, when the world state is L , the expectation of the first-round outcome is $E[\sum_i X_i \mid L] = n_F + n_C \cdot P_{hL}$. If the first round outcome is at most $x \geq n_F + n_C \cdot P_{hL} + \delta$, then all the agents i will vote for \mathbf{R} , and \mathbf{R} becomes the winner. Therefore, the probability such that \mathbf{R} becomes the winner

conditioned on the world state L satisfies

$$\begin{aligned}
\lambda_L^R(\Sigma^*) &\geq \Pr \left[\sum_i X_i \leq n_F + n_C \cdot P_{hL} + \delta \cdot n \mid L \right] \\
&= \Pr \left[\sum_i X_i - E \left[\sum_i X_i \mid L \right] \leq n_F + n_C \cdot P_{hL} + \delta \cdot n - E \left[\sum_i X_i \mid L \right] \mid L \right] \\
&= \Pr \left[\sum_i X_i - E \left[\sum_i X_i \mid L \right] \geq \delta \cdot n \mid L \right] \\
&\geq 1 - \exp(2\delta^2 \cdot n).
\end{aligned}$$

Therefore, both $\lambda_H^A(\Sigma_n^*)$ and $\lambda_L^R(\Sigma_n^*)$ converge to 1 as n goes to infinity. Since fidelity $A(\Sigma^*) = P_H \cdot \lambda_H^A(\Sigma_n^*) + P_L \cdot \lambda_L^R(\Sigma_n^*)$, $\lim_{n \rightarrow \infty} A(\Sigma^*) = 1$. \square

LEMMA 3. Σ_n^* is an ε -strong BNE with $\varepsilon = o(1)$.

By applying Lemma 4, we only need to show that the deviating group containing only one type of predetermined agents cannot succeed. Therefore, “informative + threshold” strategy profile Σ^* is an ε -strong BNE with ε converges to 0.

CLAIM 3. Suppose D is the deviating group that contains only friendly or only unfriendly agents, then $u_i(\Sigma') - u_i(\Sigma^*) < \varepsilon$ for all $i \in D$.

We prove the friendly agent side, and the unfriendly agent side will follow the same reasoning. The expected utility of a friendly agent increases only when the likelihood that A becomes the winner increases. However, in Σ_n^* , all the friendly agents always vote for A in both rounds. If friendly agents deviate in the second round, there will be strictly fewer agents voting for A in expectation, and A will be less likely to be the winner. If friendly agents deviate in the first round, the result of the first-round voting x will have a higher probability of being smaller, and contingent agents playing the threshold strategy will tend to vote R. Therefore, no matter how D deviates, the likelihood of A being the winner will not increase, and friendly agents cannot get a gain of more than ε .

B.3 Proof for Lemma 5

LEMMA 5. The sequence of “informative + sincere” strategy profiles is constantly separable.

PROOF. Recall that

$$x^* = n_F + \frac{\log \frac{P_L \cdot (v_i(L,R) - v_i(L,A))}{P_H \cdot (v_i(H,A) - v_i(H,R))} + n_C \cdot \log \frac{P_{\ell L}}{P_{\ell H}}}{\log \frac{P_{hH} P_{\ell L}}{P_{\ell H} P_{hL}}}.$$

As n_F and n_C increases with n , while $\log \frac{P_L \cdot (v_i(L,R) - v_i(L,A))}{P_H \cdot (v_i(H,A) - v_i(H,R))}$ remains constant, it suffice to show that

$$n_F + n_C \cdot P_{hL} + \delta \cdot n \leq n_F + \frac{n_C \cdot \log \frac{P_{\ell L}}{P_{\ell H}}}{\log \frac{P_{hH} P_{\ell L}}{P_{\ell H} P_{hL}}} \leq n_F + n_C \cdot P_{hH} - \delta \cdot n,$$

which is equivalent to

$$P_{hL} + \delta \cdot \frac{n}{n_C} \leq \frac{\log \frac{P_{\ell L}}{P_{\ell H}}}{\log \frac{P_{hH} P_{\ell L}}{P_{\ell H} P_{hL}}} \leq P_{hH} - \delta \cdot \frac{n}{n_C}.$$

\square

We first deal with the left inequality.

$$\begin{aligned}
P_{hL} + \delta \cdot \frac{n}{n_C} &\leq \frac{\log \frac{P_{\ell L}}{P_{\ell H}}}{\log \frac{P_{hH} P_{\ell L}}{P_{\ell H} P_{hL}}} \\
\Leftrightarrow P_{hL} + \delta \cdot \frac{n}{n_C} &\leq \frac{\log \frac{P_{\ell L}}{P_{\ell H}}}{\log \frac{P_{hH}}{P_{hL}} + \log \frac{P_{\ell L}}{P_{\ell H}}} \\
\Leftrightarrow \delta \cdot \frac{n}{n_C} \cdot \left(\log \frac{P_{hH}}{P_{hL}} + \log \frac{P_{\ell L}}{P_{\ell H}} \right) &\leq \left(P_{hL} \cdot \log \frac{P_{hL}}{P_{hH}} + \log P_{\ell L} \cdot \frac{P_{\ell L}}{P_{\ell H}} \right).
\end{aligned}$$

Note that the RHS takes a form of the KL-divergence between distribution $(P_{hL}, P_{\ell L})$ and $(P_{hH}, P_{\ell H})$. Given that $P_{hH} > P_{hL}$, RHS is a strictly positive constant independent from n . On the other hand, $n_C \approx \alpha_C \cdot n$. Therefore, by taking

$$\delta \leq \frac{1}{2} \alpha_C \cdot \frac{\left(P_{hL} \cdot \log \frac{P_{hL}}{P_{hH}} + \log P_{\ell L} \cdot \frac{P_{\ell L}}{P_{\ell H}} \right)}{\left(\log \frac{P_{hH}}{P_{hL}} + \log \frac{P_{\ell L}}{P_{\ell H}} \right)},$$

the left inequality is satisfied.

Likewise, the right inequality is equivalent to

$$\delta \cdot \frac{n}{n_C} \cdot \left(\log \frac{P_{hH}}{P_{hL}} + \log \frac{P_{\ell L}}{P_{\ell H}} \right) \leq \left(P_{hH} \cdot \log \frac{P_{hH}}{P_{hL}} + \log P_{\ell H} \cdot \frac{P_{\ell H}}{P_{\ell L}} \right),$$

and by taking

$$\delta \leq \frac{1}{2} \alpha_C \cdot \frac{\left(P_{hH} \cdot \log \frac{P_{hH}}{P_{hL}} + \log P_{\ell H} \cdot \frac{P_{\ell H}}{P_{\ell L}} \right)}{\left(\log \frac{P_{hH}}{P_{hL}} + \log \frac{P_{\ell L}}{P_{\ell H}} \right)},$$

the right inequality is satisfied. Therefore, by a sufficiently small δ and all sufficiently large n , the constant for constantly separable is satisfied.

B.4 Proof of Theorem 1

The proof applies to both the 0-1 utility setting (Theorem 1) and the general valuation utility setting (Theorem 1'). Before starting the proof, we extend the definition of the first-round voting share. let $\varphi_H^F(\Sigma)$, $\varphi_H^U(\Sigma)$, and $\varphi_H^C(\Sigma)$ be the expectation on the number of votes for **A** from three types of agents conditioned on the world state is H and the strategy profile is Σ . We have $\varphi_H(\Sigma) = \varphi_H^F(\Sigma) + \varphi_H^U(\Sigma) + \varphi_H^C(\Sigma)$. The L side is defined similarly.

PROOF. We first give the proof for regular profiles. Suppose $\{\Sigma_n\}$ is a regular profile sequence such that $\lim_{n \rightarrow \infty} A(\Sigma_n) = 1$ does not hold. This means that there exist a constant δ and an infinite set \mathcal{N} such that for all $n \in \mathcal{N}$, $A(\Sigma_n) \leq 1 - \delta$. We show that there exists a constant $\varepsilon = \frac{\delta \cdot P_H P_L}{2}$ and a infinite set \mathcal{N}' such that for all $n \in \mathcal{N}'$, Σ_n is NOT an ε -strong BNE.

We will first construct a different profile sequence $\{\Sigma'_n\}$ whose fidelity converges to 1. Then for all sufficiently large $n \in \mathcal{N}$, both λ_H^A and λ_L^R are close to one. On the other hand, Since $A(\Sigma_n) \leq 1 - \delta$, we have at least one of $\lambda_H^A(\Sigma_n) \leq 1 - \delta$ and $\lambda_L^R(\Sigma_n) \leq 1 - \delta$ holds. Then we can see that all the contingent agents have incentives to deviate to Σ'_n to get a constant gain in the expected utility.

For simplicity, we use Σ and Σ' to denote Σ_n and Σ'_n . Let Σ'_n be the following strategy profile. The deviating group D contains all the contingent agents.

- All predetermined agents play the same strategy as in Σ .
- If $\varphi_H^F(\Sigma) + \varphi_H^U(\Sigma) \geq \varphi_L^F(\Sigma) + \varphi_L^U(\Sigma)$, all contingent agents vote informatively in the first round. For the second round, all contingent agents vote for **A** if $x \geq \frac{\varphi_H(\Sigma') + \varphi_L(\Sigma')}{2}$ and vote for **R** otherwise.

- If $\varphi_H^F(\Sigma) + \varphi_H^U(\Sigma) < \varphi_L^F(\Sigma) + \varphi_L^U(\Sigma)$, contingent agents “reverse” their strategy. In the first round, a contingent agent votes for **A** if receiving l and for **R** if receiving h . And in the second round, all contingent agents vote for **A** if $x < \frac{\varphi_H(\Sigma') + \varphi_L(\Sigma')}{2}$ and vote for **R** otherwise.

The high-level idea of Σ' is similar to Σ^* in Theorem 3: contingent agents determine the world state with high probability via the first-round voting and vote for the correct alternative in the second round. Since predetermined agents behave regularly, and both α_F and α_U do not exceed 0.5, the alternative that all contingent agents vote for will be the winner. Now we show that both $\lim_{n \rightarrow \infty} \lambda_H^A(\Sigma') = 1$ and $\lim_{n \rightarrow \infty} \lambda_L^R(\Sigma') = 1$.

With loss of generality, suppose $\varphi_H^F(\Sigma) + \varphi_H^U(\Sigma) \geq \varphi_L^F(\Sigma) + \varphi_L^U(\Sigma)$. The reasoning for the other case will be similar. For an agent i , Let X_i be the random variable indicating i 's vote in the first round. $X_i = 0$ means i votes for **R**, and $X_i = 1$ means i votes for **A**. Then the result of the first round voting x can be represented as $x = \sum_i X_i$. Then the probability such that **A** becomes the winner (which is the probability that $x \geq \frac{\varphi_H(\Sigma') + \varphi_L(\Sigma')}{2}$) conditioned on the world state H can be written as

$$\lambda_H^A(\Sigma') = \Pr\left[\sum_i X_i \geq \frac{\varphi_H(\Sigma') + \varphi_L(\Sigma')}{2} \mid H\right].$$

Note that $E[\sum_i X_i \mid H] = \varphi_H(\Sigma')$.

Since contingent agent votes informatively in the first round, $\varphi_H^C(\Sigma') = n_C \cdot P_{hH}$ and $\varphi_L^C(\Sigma') = n_C \cdot P_{hL}$. Since $P_{hH} > P_{hL}$, we have $\varphi_H^C(\Sigma') - \varphi_L^C(\Sigma') = n_C \cdot (P_{hH} - P_{hL}) \geq n \cdot (\alpha_C - \frac{1}{n}) \cdot (P_{hH} - P_{hL})$. Then there exists a n_0 such that for all $n > n_0$, $\varphi_H^C(\Sigma') - \varphi_L^C(\Sigma') \geq \frac{1}{2}n \cdot \alpha_C \cdot (P_{hH} - P_{hL})$. Then since $\varphi_H^F(\Sigma) + \varphi_H^U(\Sigma) \geq \varphi_L^F(\Sigma) + \varphi_L^U(\Sigma)$, we have $n > n_0$, $\varphi_H(\Sigma') - \varphi_L(\Sigma') \geq \frac{1}{2}n \cdot \alpha_C \cdot (P_{hH} - P_{hL})$. Then we can apply the Hoeffding inequality to the probability:

$$\begin{aligned} \lambda_H^A(\Sigma') &= \Pr\left[\sum_i X_i \geq \frac{\varphi_H(\Sigma') + \varphi_L(\Sigma')}{2} \mid H\right] \\ &= \Pr\left[\sum_i X_i - E\left[\sum_i X_i \mid H\right] \geq \frac{\varphi_H(\Sigma') + \varphi_L(\Sigma')}{2} - E\left[\sum_i X_i \mid H\right] \mid H\right] \\ &= \Pr\left[\sum_i X_i - \varphi_H(\Sigma') \geq \frac{\varphi_L(\Sigma') - \varphi_H(\Sigma')}{2} \mid H\right] \\ &\geq \Pr\left[\sum_i X_i - \varphi_H(\Sigma') \geq -\frac{1}{4}n \cdot \alpha_C \cdot (P_{hH} - P_{hL}) \mid H\right] \\ &\geq 1 - \exp\left(-\frac{1}{8}(\alpha_C \cdot (P_{hH} - P_{hL})^2 \cdot n)\right). \end{aligned}$$

The L side works similarly.

Therefore, both $\lambda_H^A(\Sigma'_n)$ and $\lambda_L^R(\Sigma'_n)$ converge to 1 as n goes to infinity.

Next, we show that contingent agents have incentives to deviate from Σ_n to Σ'_n . Let \mathcal{N}' be the set of all $n \in \mathcal{N}$ such that both $\lambda_H^A(\Sigma'_n) \geq 1 - \frac{\delta \cdot P_H P_L}{2B}$ and $\lambda_L^R(\Sigma'_n) \geq 1 - \frac{\delta \cdot P_H P_L}{2B}$ holds. We know that \mathcal{N}' is also an infinite set. We fix an arbitrary $n \in \mathcal{N}'$. Without loss of generality, suppose $\lambda_H^A(\Sigma_n) \leq 1 - \delta$. Then we consider the difference in the expected utility of a contingent agent.

$$\begin{aligned} u_i(\Sigma'_n) - u_i(\Sigma_n) &= P_H \cdot (\lambda_H^A(\Sigma'_n) - \lambda_H^A(\Sigma_n))(v_i(H, \mathbf{A}) - v_i(H, \mathbf{R})) \\ &\quad + P_L \cdot (\lambda_L^R(\Sigma'_n) - \lambda_L^R(\Sigma_n))(v_i(L, \mathbf{R}) - v_i(L, \mathbf{A})). \end{aligned}$$

By assigning the bounds above, we have

$$\begin{aligned} u_i(\Sigma'_n) - u_i(\Sigma_n) &\geq P_H \cdot \left(1 - \frac{\delta \cdot P_H P_L}{2B} - 1 + \delta\right) \cdot 1 + P_L \cdot \left(-\frac{\delta \cdot P_H P_L}{2B}\right) \cdot B \\ &\geq P_H \cdot \delta - \frac{\delta \cdot P_H P_L}{2} \\ &> \frac{\delta \cdot P_H P_L}{2}. \end{aligned}$$

The similar reasoning can be achieved for $\lambda_L^R(\Sigma_n) \leq 1 - \delta$.

Therefore, all the contingent agents have incentives to deviate from Σ_n to Σ'_n to get a constant gain. This means that Σ_n is NOT an ε -strong BNE for some constant ε .

B.4.1 Proof for non-regular profiles. We adopt the same idea to construct a strategy profile where the deviating group forces the state to be revealed by the first-round vote and votes for the preferred agent in the second round. Then we show that as the fidelity converges to 1, the group has incentives to do so.

For non-regular profiles, a deviating group with only contingent agents may not work, as there is no assumption on the behavior of predetermined agents in the second round. (For example, if $\alpha_C < 0.5$, and all the predetermined agents always vote for A, then the winner will always be A regardless of the votes from contingent agents.) Fortunately, the following lemma shows that if the fidelity of deviating strategy profile is close to 1, then either all the friendly agents get an expected utility increase or all the unfriendly agents get an expected utility increase. Moreover, which type gets utility gain does not depend on the deviating strategy Σ'_n .

LEMMA 1. *Let $\{\Sigma_n\}$ be a sequence of profiles such that there exists a constant δ and an infinite set \mathcal{N} such that for all $n \in \mathcal{N}$, $A(\Sigma_n) \leq 1 - \delta$. Then there exists an infinite set $\mathcal{N}' \subseteq \mathcal{N}$ such that for each $n \in \mathcal{N}'$, at least one of the following holds for all sequences of profiles $\{\Sigma'_n\}$ with $\lim_{n \rightarrow \infty} A(\Sigma') = 1$:*

- (1) *For any friendly agent i_1 , $u_{i_1}(\Sigma'_n) - u_{i_1}(\Sigma_n) \geq 0$.*
- (2) *For any unfriendly agent i_2 , $u_{i_2}(\Sigma'_n) - u_{i_2}(\Sigma_n) \geq 0$.*

Moreover, for a fixed n and Σ_n , either 1 holds for all Σ'_n or 2 holds for all Σ'_n .

The proof of Lemma 1 comes from the fact that a predetermined agent has different sensitivities to utility changes in different states of the world. For a friendly agent i_1 , $v_{i_1}(H, \mathbf{A}) > v_{i_1}(L, \mathbf{A}) > v_{i_1}(L, \mathbf{R}) > v_{i_1}(H, \mathbf{R})$. Therefore, i_1 is more sensitive to utility change in world state H than that in world state L . On the other hand, an unfriendly agent i_2 is more sensitive in world state L . This difference in sensitivity guarantees that when the fidelity increases, the expected utility of both types of agents cannot decrease at the same time.

PROOF OF LEMMA 1. Without loss of generality, suppose $u_{i_1}(\Sigma'_n) - u_{i_1}(\Sigma_n) \leq 0$ holds for i_1 (which means (1) does not hold). We show that (2) $u_{i_2}(\Sigma'_n) - u_{i_2}(\Sigma_n) > 0$ must hold for i_2 . Recall that the expected utility of an agent i is defined as follows.

$$\begin{aligned} u_i(\Sigma) &= P_L(\lambda_L^A(\Sigma) \cdot v_i(L, \mathbf{A}) + \lambda_L^R(\Sigma) \cdot v_i(L, \mathbf{R})) \\ &\quad + P_H(\lambda_H^A(\Sigma) \cdot v_i(H, \mathbf{A}) + \lambda_H^R(\Sigma) \cdot v_i(H, \mathbf{R})). \end{aligned}$$

Therefore, the difference between $u_i(\Sigma'_n)$ and $u_i(\Sigma_n)$ can be represented as follows.

$$\begin{aligned} &P_H \cdot (\lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*)) (v_i(H, \mathbf{A}) - v_i(H, \mathbf{R})) \\ &+ P_L \cdot (\lambda_L^R(\Sigma') - \lambda_L^R(\Sigma^*)) (v_i(L, \mathbf{R}) - v_i(L, \mathbf{A})). \end{aligned}$$

Since $A(\Sigma_n) \leq 1 - \delta$, at least one of $\lambda_H^A(\Sigma_n) \leq 1 - \delta$ and $\lambda_L^R(\Sigma_n) \leq 1 - \delta$ holds. We deal with two cases separately.

Case 1: $\lambda_H^A(\Sigma_n) \leq 1 - \delta$. In this case, we have $\lambda_H^A(\Sigma'_n) - \lambda_H^A(\Sigma_n) > 0$. Note that for the friendly agent i_1 , $v_{i_1}(H, \mathbf{A}) - v_{i_1}(H, \mathbf{R}) \geq -(v_{i_1}(L, \mathbf{R}) - v_{i_1}(L, \mathbf{A}))$. Therefore, $u_{i_1}(\Sigma') - u_{i_1}(\Sigma) \leq 0$ implies

$$P_H \cdot (\lambda_H^A(\Sigma'_n) - \lambda_H^A(\Sigma_n)) \leq P_L \cdot (\lambda_L^R(\Sigma'_n) - \lambda_L^R(\Sigma_n)).$$

On the other hand, for any unfriendly agent i_2 , $-v_{i_2}(H, \mathbf{A}) - v_{i_2}(H, \mathbf{R}) \leq v_{i_2}(L, \mathbf{R}) - v_{i_2}(L, \mathbf{A})$. Therefore, $u_{i_2}(\Sigma'_n) - u_{i_2}(\Sigma_n) \geq 0$.

Case 2: $\lambda_L^R(\Sigma_n) \leq 1 - \delta$. In this case, we have $\lambda_L^R(\Sigma'_n) - \lambda_L^R(\Sigma_n) > 0$. If $\lambda_H^A(\Sigma'_n) - \lambda_H^A(\Sigma_n) \leq 0$, then

$$u_{i_2}(\Sigma'_n) - u_{i_2}(\Sigma_n) \geq P_L \cdot (\lambda_L^R(\Sigma'_n) - \lambda_L^R(\Sigma_n))(v_i(L, \mathbf{R}) - v_i(L, \mathbf{A})) > 0.$$

If $\lambda_H^A(\Sigma'_n) - \lambda_H^A(\Sigma_n) > 0$, we apply the same reasoning as in Case 1. Therefore, $u_{i_2}(\Sigma'_n) - u_{i_2}(\Sigma_n) \geq 0$ always holds. For the 0-1 utility model, all these differences utilities become 1 or -1 , and suffice to compare $P_H \cdot (\lambda_H^A(\Sigma'_n) - \lambda_H^A(\Sigma_n))$ and $P_L \cdot (\lambda_L^R(\Sigma'_n) - \lambda_L^R(\Sigma_n))$ to determined which type of agents is happy. \square

Now we construct the deviating strategy profile sequence $\{\Sigma'_n\}$, which is a combination of two strategy profile sequences $\{\Sigma_F\}$ and $\{\Sigma_U\}$. In both Σ_F and Σ_U contingent agents collaborate with one type of predetermined agents and deviate together.

In Σ_F , all the friendly and contingent agents vote informatively in the first round and play a threshold strategy in the second round.

- All *unfriendly* agents play the same strategy as in Σ .
- If $\varphi_H^U(\Sigma_n) \geq \varphi_L^U(\Sigma_n)$, all friendly agents and contingent agents vote informatively in the first round. For the second round, all friendly agents and contingent agents vote for \mathbf{A} if $x \geq \frac{\varphi_H(\Sigma_F) + \varphi_L(\Sigma_F)}{2}$ and vote for \mathbf{R} otherwise.
- If $\varphi_H^U(\Sigma_n) < \varphi_L^U(\Sigma_n)$, friendly agents and contingent agents “reverse” their strategy. In the first round, a friendly (or contingent) agent votes for \mathbf{A} if receiving l and for \mathbf{R} if receiving h . In the second round, all friendly agents and contingent agents vote for \mathbf{A} if $x \geq \frac{\varphi_H(\Sigma_F) + \varphi_L(\Sigma_F)}{2}$

Σ_U follows the same “informative and threshold” pattern, except that the deviators are all the unfriendly agents and all the contingent agents.

Since $\alpha_F < 0.5$ and $\alpha_U < 0.5$, if contingent agents and one type of predetermined agents vote for the same alternative, that alternative will be the winner. Therefore, with similar reasoning to the regular case, both Σ_F and Σ_U have fidelity converging to 1, and contingent agents have incentives to deviate to either of them.

Applying Lemma 1, we know that for each n where the fidelity of Σ_n does not converge to 1, there is one type of predetermined agents whose expected utilities will increase if the deviators deviate to either Σ_F or Σ_U . Moreover, the type of agents that get higher expected utilities does not depend on which deviating strategy they play. Therefore, when friendly agents get higher expected utilities, the deviators are all the friendly agents and contingent agents with $\Sigma'_n = \Sigma_F$; otherwise, the deviators are all the unfriendly agents and contingent agents with $\Sigma'_n = \Sigma_U$. Then, the expected utilities of contingent increase by at least a constant, and those of the deviating predetermined agents increase. Therefore, for each Σ_n whose fidelity is not close to 1, the deviators have incentives to deviate to either Σ_F or Σ_U . Therefore, Σ_n is NOT an ε -strong BNE with constant ε . \square

B.5 Formal Statement and Proof of Theorem 4

THEOREM 4'. *Let $\{\hat{\Sigma}_n\}$ be a one-round regular strategy profile sequence in an environment sequence $\{\mathcal{I}_n\}$. If for all n , $\hat{\Sigma}$ is an $\bar{\varepsilon}$ -strong BNE in one-round voting with $\bar{\varepsilon}$ converges to 0, then for any two-round*

strategy profile sequence $\{\Sigma_n\}$ in environment $\{\mathcal{I}_n\}$, where for any n , Σ_n is a second-round consistent strategy profile of $\hat{\Sigma}_n$, Σ_n is an ε -strong BNE in the two-round voting for all n with ε converges to 0.

PROOF. The proof proceeds in two steps. Firstly, we show that the fidelity of $\{\Sigma_n\}$ converges to 1. Secondly, by applying Lemma 4, we know that a deviating group either contains only friendly agents or only unfriendly agents. Finally, we show that such a deviating group cannot succeed.

Step 1: $A(\Sigma_n)$ converges to 1. In Σ_n , the behavior and result in the first-round voting have no effect to the second-round voting that determines the winner. Agents in the second round behave exactly as they are playing a one-round voting game. Therefore, $A(\Sigma_n) = A(\hat{\Sigma}_n)$. Then since $\hat{\Sigma}_n$ is an $\bar{\varepsilon}$ -strong BNE with $\bar{\varepsilon}$ converges to 0, we have $A(\hat{\Sigma}_n)$ converges to 1 according to the if and only if characterization in Han et al. [2023, Theorem 2]. Therefore, $A(\Sigma_n)$ also converges to 1.

Step 2: equilibrium. For each n , let $\varepsilon = 2B(B+2)(1 - A(\Sigma_n))$. Fix an arbitrary n . Suppose there exists a strategy profile Σ'_n such that a deviating group D of agents have incentives to deviate to Σ'_n , then by Lemma 4, D either contains only friendly agents or only unfriendly agents.

Then we show that such D cannot succeed. This follows the last part of Lemma 3. With loss of generality, suppose D contains only friendly agents. Since $\hat{\Sigma}_n$ is a regular profile, we know that in Σ_n , friendly agents always vote for A in the second round. Then friendly agents cannot make themselves better off. The changes in the first-round voting have no effect on the second round, and changes in the second-round voting will only decrease the votes for A and cannot increase the probability that A becomes the winner. Therefore, for any friendly agents, the expected utility cannot increase by ε .

Therefore, we have show that for each n , Σ_n is a ε -strong BNE with $\varepsilon = 2B(B+2)(1 - A(\Sigma_n))$. Since the fidelity converges to 1, ε will converge to 0. \square

C More Details on Experiments.

C.1 Configuration

The experiment is run on a Lenovo Thinkbook 15 G3 AGL with AMD Ryzen 5 5600U with Radeon Graphics 2.30 GHz. The program is written in Python 3.9.

We adopt online APIs to create generative AI agents. The API for Deepseek R1 is provided by Together AI, and the API for GPT-4o is provided by OpenAI's official API platform.

C.2 Examples for prompts and answers of agents.

Here we give two examples of the prompts and the answers of the agents. More data can be found in the supplementary materials.

C.2.1 Deepseek R1, biased signal, accurate information, and one-round voting.

Signal. Bad Impression.

Prompt. "Consider the following hypothetical situation. Suppose you are a recruiter in a company. Your company is hiring for a new position by voting. There are in total 40 voters. You believe that the candidate should be hired only if he is qualified. Most of the company agrees with you, but a small minority <about 12%> seem to support hiring every candidate regardless of qualifications, and another small minority <about 7%> does not support hiring any candidate regardless of qualifications. Before meeting with the candidate, you thought she has <about 50% chance> to be qualified. If the candidate is qualified, you will likely have a good impression after meeting with her <about 80% chance>. If the candidate is not qualified, you will be less likely to have a good impression <about 60% chance>. It seems that everyone is on the same page about both the candidate's likelihood of being qualified and how informative an interview is, but the interview impression is private and independent among all the voters. You got a BAD impression after meeting

the candidate. After interviewing, but before any discussions, the company holds a vote. There will be only one round, where each voter vote YES or NO on whether to hire this candidate. The outcome is determined by whether the candidate gets a majority of yes votes. Analyze what is your best vote. The voting time is limited, so you cannot think everlastingly. Summarize and make an immediate decision in two minutes. Your answer should start with the analysis and end with a single character Y or N in a separate line without any other characters or symbols, indicating your vote. ",

Vote. N.

Answer. The candidate has a 33.3% posterior probability of being qualified after a bad impression. With 5 automatic YES and 3 NO votes, the remaining 32 voters' expected YES votes (weighted by qualification likelihood) sum to $21.33 + 5 = 26.33$, likely securing a majority. However, given the higher probability the candidate is unqualified (66.7%), voting NO aligns with preventing unqualified hires despite the likely majority.

C.2.2 Deepseek R1, biased signal, vague information, and two-round voting. Here we give an example where an agent in two-round voting switches its vote in the second round.

Signal. Good Impression.

First-round Prompt. Consider the following hypothetical situation. Suppose you are a recruiter in a company. Your company is hiring for a new position by voting. There are in total 40 voters. You believe that the candidate should be hired only if he is qualified. Most of the company agrees with you, but a small minority seem to support hiring every candidate regardless of qualifications, and another small minority does not support hiring any candidate regardless of qualifications. Before meeting with the candidate, you thought her likelihood of being qualified was roughly balanced. If the candidate is qualified, you would almost certainly form a favorable impression after meeting her. If the candidate is not qualified, you would be far less inclined to have a positive impression, with odds comparable to a rare occurrence. It seems that everyone is on the same page about both the candidate's likelihood of being qualified and how informative an interview is, but the interview impression is private and independent among all the voters. You got a GOOD impression after meeting the candidate. After interviewing, but before any discussions, the company holds a vote. There will be two rounds. In each round a voter votes yes or no. In the first round, the individual votes are private, but the number of votes will be announced to all the voters. This round does not determine the outcome. In the second round, the voting is also private., the outcome is determined by whether the candidate gets a majority of YES votes. Analyze what is your best vote in the first round. The voting time is limited, so you cannot think everlastingly. Summarize and make an immediate decision in two minutes. Your answer should start with the analysis and end with a single character Y or N in a separate line without any other characters or symbols, indicating your vote.

First-round Vote. Y.

First-round Answer. To determine the optimal first-round vote, consider Bayesian updating: a GOOD impression strongly indicates the candidate is qualified (91% probability). The first-round vote count informs the second-round voters. Voting Y truthfully contributes to an accurate signal, aiding rational voters in updating their beliefs. Extremists' votes are fixed, but their impact is diluted if rational voters coordinate via the first-round tally. Truthful voting maximizes collective accuracy.

Second-round Prompt. Consider the following hypothetical situation. Suppose you are a recruiter in a company. Your company is hiring for a new position by voting. There are in total 40 voters. You believe that the candidate should be hired only if he is qualified. Most of the company agrees with you, but a small minority seem to support hiring every candidate regardless of qualifications, and another small minority does not support hiring any candidate regardless of qualifications. Before meeting with the candidate, you thought her likelihood of being qualified was roughly balanced. If the candidate is qualified, you would almost certainly form a favorable impression after meeting her. If the candidate is not qualified, you would be far less inclined to have a positive impression, with odds comparable to a rare occurrence. It seems that everyone is on the same page about both the candidate's likelihood of being qualified and how informative an interview is, but the interview impression is private and independent among all the voters. You got a GOOD impression after meeting the candidate. After interviewing, but before any discussions, the company holds a vote. There will be two rounds. In each round, a voter votes yes or no. In the first round, the individual votes are private, but the number of votes will be announced to all the voters. This round does not determine the outcome. In the second round, the voting is also private., the outcome is determined by whether the candidate gets a majority of YES votes. After the first round, 14 of 40 voters voted yes. Analyze what is your best vote in the second round. The voting time is limited, so you cannot think everlastingly. Summarize and make an immediate decision in two minutes. Your answer should start with the analysis and end with a single character Y or N in a separate line without any other characters or symbols, indicating your vote.

Second-round Vote. N.

Second-round Answer. After observing 14/40 first-round YES votes, Bayesian updating suggests the candidate is likely unqualified. Given rational voters' independent impressions, a low YES count aligns more with an unqualified candidate (expected 8-12 YES) than qualified ($\tilde{28}$ -32). Despite my GOOD impression, the aggregate signal outweighs individual likelihood due to extreme unlikelihood of 14 YES under qualification. Vote NO to align with probable unqualification.