

---

# Wisdom of the Crowd Voting: Truthful Aggregation of Voter Information and Preferences

---

**Grant Schoenebeck**  
School of Information  
University of Michigan  
schoeneb@umich.edu

**Biaoshuai Tao\***  
John Hopcroft Center for Computer Science  
Shanghai Jiao Tong University  
bstao@sjtu.edu.cn

## Abstract

We consider two-alternative elections where voters' preferences depend on a state variable that is not directly observable. Each voter receives a private signal that is correlated to the state variable. Voters may be "contingent" with different preferences in different states; or predetermined with the same preference in every state. In this setting, even if every voter is a contingent voter, agents voting according to their private information need not result in the adoption of the universally preferred alternative, because the signals can be systematically biased.

We present an easy-to-deploy mechanism that elicits and aggregates the private signals from the voters, and outputs the alternative that is favored by the majority. In particular, voters truthfully reporting their signals forms a strong Bayes Nash equilibrium (where no coalition of voters can deviate and receive a better outcome).

## 1 Introduction

Social choice theory studies how to aggregate participants' heterogeneous opinions/preferences and output a collective decision from a set of alternatives. Typically, though not always, it is assumed that each participant has a clear preference over the alternatives, e.g., a preference order over all the alternatives, a valuation for each alternative, etc. However, even with only two alternatives, this is not the typical case for *all* participants. In addition to the participants who have clear, *predetermined* preferences for one alternative over the other, typically there are also *contingent* participants who only have partial information on which alternative is "preferable for them" and yet would like to select the alternative that is "preferable for them."

A standard example would be an election with two candidates  $a$  and  $b$  coming from political parties  $A$  and  $B$  respectively. Voters are normally partitioned into three types: some partisans for Party  $A$  prefer candidate  $a$  based on his support for the platform of party  $A$ ; other partisans for Party  $B$  prefer candidate  $b$  based on her support for the platform of party  $B$ ; finally, there are swing voters who are largely indifferent between the parties' platforms and would like to elect whichever candidate can make more progress on non-partisan issues. However, swing voters do not have perfect information about which candidate is better suited for addressing the non-partisan needs of the community. Instead, each voter has a hunch of which candidate will perform better on the non-partisan problems facing the community based on both public information and their private experiences and beliefs.

Additional examples where participants have preferences, but may or may not know what is "preferable for them", abound. In votes for corporate strategies, for hiring decisions, and for policy decisions typically some participants would like to select according to some truth they are collectively trying to

---

A full version of this paper is available at <https://arxiv.org/abs/2108.03749>.

This material is based upon work supported by the National Science Foundation under Grant No. 2007256.

\*corresponding author

discern (e.g., impact on future profits, suitability for the position, efficacy of policy, etc.) while others may have predetermined preferences, for example, because of the way they are uniquely affected (e.g., prominence of their position in future corporate strategy, vision/skills of the job candidate, who in particular the policy benefits/harms).

## 1.1 Informal Setting

In this paper, we consider a two-alternative social choice setting where voters' preferences may depend on a state variable. In the above example, the state is which candidate will make more progress on non-partisan issues.

We consider binary state variables in the main body of this paper. In general, the state need not be binary. For example, we could generalize the above example so which candidate would perform better on non-partisan issues is on a scale of 1 to 10 (where 1 indicates candidate  $a$  is much better and 10 indicates candidate  $b$  is much better). Predetermined voters' preferences would still not depend on the state. However, contingent voters' may have different thresholds on the state where they would transfer their support from candidate  $a$  to candidate  $b$ . In Appendix E, we discuss the non-binary setting and extend our results to this setting.

The state variable is not directly observable in the election phase, as the performance of a new government official, a new hired employee, a new policy, etc., may not be revealed until many years after the vote. Instead, each voter receives a signal that is correlated to the hidden state which models the information voters received from different sources.

Our goal is to select the *majority wish*, the alternative that would be preferred by the majority if they knew the state of the world. If some type of predetermined voters forms a majority, this is rather easy. However, in the case where the predetermined voters of neither alternative forms a majority, the mechanism needs to aggregate the information and preferences of the contingent voters and selects the alternative which, for a majority of voters, is "preferable for them".

## 1.2 Imperfectly Informed Voters

**Social Choice** Social choice theory with imperfectly informed voters dates back to Condorcet's jury theorem in 1785 [De Condorcet, 1785], and has also been widely studied [Miller, 1986, Young, 1988, Ladha, 1992, Nitzan and Paroush, 2017]. In Condorcet's setting, there are two alternatives, one of which is "correct", and each voter votes for the correct alternative with probability  $p$ . Condorcet's jury theorem states that the probability that the majority voting scheme outputs the correct alternative goes to 1 as the number of voters increases when  $p > 0.5$ , and, conversely, this probability goes to 0 when  $p < 0.5$ .

Two unfortunate limitations for Condorcet's jury theorem are 1) It fails to output the correct alternative in the case voters' beliefs are aligned to the incorrect alternative (i.e.,  $p < 0.5$ ), and 2) It assumes voters vote truthfully and disregards voters' potential strategic behaviors. However, even in the case all voters have the same preference for the correct alternative, voting truthfully still may not be a Nash equilibrium [Austen-Smith and Banks, 1996].

To circumvent the first limitation, Feddersen and Pesendorfer [1997] consider the scenario where voters play a Nash equilibrium strategy profile, while the voting rule is still the majority scheme. Feddersen and Pesendorfer [1997] show that when the number of voters tends to infinity, the probability that a majority voting scheme outputs the correct alternative approaches to 1 if voters play the equilibrium strategy profile, while this probability is bounded away from 1 if voters play the truthful strategy profile instead. Feddersen and Pesendorfer's model assigns each voter a preference parameter  $x \in [-1, 1]$  describing his/her alignment to the two alternatives. In their model, the unique (Bayes) Nash equilibrium is characterized by two thresholds  $x_0, x_1 \in [-1, 1]$  with  $x_0 < x_1$  such that voters with preferences below  $x_0$  always vote for one alternative, voters with preferences above  $x_1$  always vote for the other, and voters with preferences between  $x_0$  and  $x_1$  vote truthfully. Although Feddersen and Pesendorfer's solution guarantees that the correct alternative is output with high probability, it requires sophisticated voters. The voters need to calculate the values of  $x_0$  and  $x_1$  to decide their actions. The values  $x_0$  and  $x_1$  are each the zero point of a continuous monotone function involving a complicated Riemann integral. This is usually too demanding for voters in practice,

especially those who do not have a mathematical background. Moreover, it needs to be common knowledge that all agents can and will perform this computation.

In this paper, we take a different approach. Instead of asking voters to play the Nash equilibrium for majority voting, we seek to design a more sophisticated voting scheme, or a mechanism, than the majority voting scheme, such that voters are incentivized to vote *truthfully* under the mechanism, while guaranteeing the correct alternative is output with high probability. Our social choice mechanism thus elicits truthful information from the voters and then aggregates it.

**Information Aggregation** The information aggregation literature considers how to obtain a “correct answer” by aggregating individuals’ partial information—the crowd’s wisdom. The straightforward procedure of outputting the answer that is believed to be correct by the majority does not always work [Chen et al., 2004, Simmons et al., 2011]. An example where this fails is when the crowd has a strong prior belief for the incorrect answer while novel specialized knowledge is only shared among a minority of the agents. It is also known that further calibration based on collecting participants’ confidences (the posterior of their beliefs) does not always solve this problem [Hertwig, 2012, Prelec et al., 2017]. In a seminal work by Prelec et al. [2017], a new “surprisingly popular” approach was proposed: the participants’ predictions over the other remaining participants’ reported answers are collected, and the answer that is reported by more participants than predicted is output (we will review this in Section 3.1). They justified this approach both theoretically and through experimentation. In particular, they demonstrate the viability of approaches that require agents to predict reports of the other agents. Hosseini et al. [2021] empirically extend the surprisingly popular approach to the non-binary setting, where the goal is to learn the correct *ranking over many options* instead of the correct answer in two options.

The work by Prelec et al. [2017] and Hosseini et al. [2021] does not fit into the social choice context in two aspects. Firstly, the objective for an information aggregation mechanism is to output the correct answer. Participants who collaboratively contribute their knowledge/information do not have preferences on which answer is finally selected. This is fundamentally different from the social choice setting where the whole point of a social choice mechanism is to select an alternative favored by the majority. Secondly, as agents care about the outcome, agents may be strategic and manipulate their reports in order to make their preferred alternatives win, while Prelec et al. [2017] and Hosseini et al. [2021] do not put the problem in a game theory setting.

### 1.3 Our Results

In this paper, we study the social choice problem in a game theory setting with the existence of imperfectly informed voters who only have partial information regarding which alternative is more favorable. For various settings with two alternatives, we propose a mechanism that aggregates participants’ private information and outputs the alternative favored by more than half of the participants—the *majority wish*. Our mechanisms are truthful, in the sense that the truthful strategy profile forms a *strong Bayes Nash Equilibrium*.

Our main wisdom-of-the-crowd-voting mechanism, presented in the paper, applies to the case of two worlds/states, where each agent receives a binary signal. We show it has strong truthfulness and aggregation properties, even for relatively small numbers of agents (Section 3). This result requires that the distribution of agents is a common knowledge. In the full version, we show several additional results. First, the common knowledge assumption is necessary to attain a strongly truthful mechanism that outputs the majority wish with high probability (Appendix D). Additionally, we extend our results to the setting of more than two worlds (Appendix E.2). Finally, we show how to extend our results to the case of more than two signals (Appendix E.3).

Our mechanism can easily be implemented using a simple questionnaire that elicits voters’ information and preferences. The questions in the questionnaire are friendly to those voters who do not have relevant backgrounds in mathematics, game theory, etc., and they require participants to predict the responses of other agents, which is empirically validated by the surprisingly popular method from Prelec et al. [2017].

We ensure our mechanisms have a group truthfulness property by employing a “median trick.” Intuitively, by the Median Voter Theorem [Black, 1948, Hotelling, 1929], the median voter’s vote (in a binary choice) is favored by the majority. By a careful design, our mechanism ensures that

the voters who are “below” the median have a conflict of interest to the voters who are “above” the median, which makes sure less than half of the voters have an incentive to deviate and those voters can only change the outcome in the unfavorable direction by the property of median.

From a high level, our work can be understood as a revelation principal applied to plurality voting.<sup>4</sup> However this view is not entirely accurate. First, our equilibrium concept is strong Bayes Nash equilibrium while we only know that plurality voting implements the majority wish outcome in (Bayesian) Nash equilibrium [Feddersen and Pesendorfer, 1997]. Second, the revelation principal requires that agents report all their knowledge. In our case, this would include the entire prior, which is not realistic. In contrast, our mechanisms only require that agents report a preference and a prediction of other agents’ preferences. Such reporting requirements have previously been shown to be pragmatic [John et al., 2012, Prelec et al., 2017, Hosseini et al., 2021]. Third, our setting is different than prior work [Feddersen and Pesendorfer, 1997], and this makes our results incomparable. In particular, we deal with a discrete state space. This difference also allows us to achieve some of our results not just in the limit, but for finite sets of agents.

Appendix A contains an additional comparison of our work with Feddersen and Pesendorfer [1997]’s.

## 1.4 Additional Related Work

Our work is additionally related to the recent work on incentive compatible machine learning [Perote and Perote-Pena, 2004, Dekel et al., 2010, Chen et al., 2018]. In these settings, the “social choice” being made is a machine learning predictor where agents benefit from their point having small error with respect to the chosen predictor. As in our setting, the information of the optimal model is distributed among the agents. Unlike our model, the private information and the preferences of the agents essentially coincide.

Information elicitation without verification, sometimes call peer prediction, is another very related line of research which shares some of the intuitions and techniques from information aggregation. The information elicitation literature has been well established in the past decades, starting from Prelec [2004]’s Bayesian Truth Serum and Miller et al. [2005]’s peer-prediction method. These mechanisms cleverly design payments to the agents to guarantee that the truthful reporting of received information forms a Nash equilibrium.

A mass of recent work (see Faltings and Radanovic [2017] for a survey) is dedicated to designing information elicitation mechanisms that work in more general settings (such as, supporting a small number of agents [Dasgupta and Ghosh, 2013, Zhang and Chen, 2014, Schoenebeck and Yu, 2020], allowing agents having information with different levels of sophistication [Gao et al., 2016, Kong and Schoenebeck, 2018b]), or achieving better truthful guarantees (such as, strict Nash equilibrium [Schoenebeck and Yu, 2021], informed Nash equilibrium [Shnayder et al., 2016], or even dominant strategy equilibrium [Kong and Schoenebeck, 2019, Kong, 2020]) sometimes by studying more restrictive settings (e.g. multiple similar questions being asked simultaneously [Dasgupta and Ghosh, 2013]). Indeed, following Bayesian Truth Serum, many of these mechanism require the agents to predict other agents’ reports [Zhang and Chen, 2014, Kong and Schoenebeck, 2018a,b, 2019, Kong et al., 2020, Schoenebeck and Yu, 2020].

However, all these mechanisms rely on *payments* to the agents to incentivize truth-telling. In our social choice setting, on the other hand, we need to incentivize truth-telling solely based on choosing the winning alternative.

## 2 Model and Preliminaries

In this paper, we will define our model and present our main result with two states and two signals. The extension to general numbers of states and signals is discussed in Appendix E. We use the standard signal structure that is commonly used in the peer prediction (e.g., the work in Sect. 1.4), information aggregation (e.g., [Prelec et al., 2017]), and social choice literature (e.g., [Feddersen and Pesendorfer, 1997]).

---

<sup>4</sup>Loosely speaking, the revelation principal states that any outcome that can be implemented in equilibrium can also be truthfully implemented in equilibrium by having the mechanism play the equilibrium strategy on behalf of the truthful agents.

Suppose a department of  $T$  faculty members, or agents, need to decide whether or not to hire a new faculty candidate. In our model, those  $T$  agents are voting for two *alternatives*,  $\mathbf{A}$  and  $\mathbf{R}$  (corresponding to “accept” and “reject”). There is a set of 2 possible *worlds* (or *states*)  $\mathcal{W} = \{L, H\}$ , which describes the underlying quality of the candidate. Here,  $L$  stands for “low quality”, and  $H$  stands for “high quality”. Agents do not know which world is the actual world that they are in. They have a common prior belief on the likelihood of each world. In the candidate hiring example, the CV of the candidate is given to those  $T$  faculty members before any individual interviews, and a prior belief is formed. Let  $W$  be the actual world which is viewed as a random variable. We use  $W = L$  and  $W = H$  to denote the events that the actual world is  $L$  and  $H$  respectively. Let  $(P_L, P_H) = (\Pr(W = L), \Pr(W = H))$  be the prior over worlds. Each agent knows the values of  $P_L$  and  $P_H$  as prior beliefs. We further assume  $P_L, P_H > 0$ .

An individual interview for this candidate is held for each of the  $T$  agents. Each agent  $t$  receives a *signal*, represented by the random variable  $S_t$ , from the set  $\mathcal{S} = \{\ell, h\}$ . Given  $W = L$  or  $W = H$ , the signals agents receive have the same distribution and are conditionally independent. Let  $P_{\ell L} = \Pr(S_t = \ell \mid W = L)$  be the probability that signal  $\ell$  will be received (by an arbitrary agent  $t$ ) if the actual world is  $L$ . Let  $P_{hL}, P_{\ell H}$  and  $P_{hH}$  have similar meanings. The set of values  $\{P_{\ell L}, P_{\ell H}, P_{hL}, P_{hH}\}$  is known by all the agents, but not the mechanism (such mechanisms are called *detail-free* in the peer prediction literature). Naturally, signals are positively correlated to the worlds:

$$P_{\ell L} > P_{\ell H} \quad \text{and} \quad P_{hH} > P_{hL}. \quad (1)$$

Different agents may have different bars on the quality of the candidate. For example, if a theory candidate is interviewed in a computer science department, the theory faculty members may want to accept this candidate at a lower quality compared with the AI, software and hardware faculty members. To capture this, each agent  $t$  is assigned a *utility function*  $v_t : \mathcal{W} \times \{\mathbf{A}, \mathbf{R}\} \rightarrow \{0, 1, \dots, B\}$ . Naturally, voters receive higher utilities for  $\mathbf{A}$  in world  $H$  and for  $\mathbf{R}$  in world  $L$ :

$$v_t(H, \mathbf{A}) > v_t(L, \mathbf{A}) \quad \text{and} \quad v_t(H, \mathbf{R}) < v_t(L, \mathbf{R}). \quad (2)$$

Since we can always rescale agents’ utilities, for simplicity, we assume without loss of generality that agents’ utilities are integers and bounded by  $B \in \mathbb{Z}^+$ . Endowed with their prior beliefs, upon receiving their signals, agents will have posterior beliefs about the distribution of  $W$  and react to the mechanism in a way maximizing their expected utilities accordingly.

We assume  $v_t(L, \mathbf{A}) \neq v_t(L, \mathbf{R})$  and  $v_t(H, \mathbf{A}) \neq v_t(H, \mathbf{R})$  for each agent  $t$ , so that agents always strictly prefer one alternative over the other. Let  $F$  be the set of the *candidate-friendly* agents  $t$  who always prefer  $\mathbf{A}$ :  $v_t(H, \mathbf{A}) > v_t(L, \mathbf{A}) > v_t(L, \mathbf{R}) > v_t(H, \mathbf{R})$ . Let  $U$  be the set of the *candidate-unfriendly* agents  $t$  who always prefer  $\mathbf{R}$ :  $v_t(L, \mathbf{R}) > v_t(H, \mathbf{R}) > v_t(H, \mathbf{A}) > v_t(L, \mathbf{A})$ . Let  $C$  be the set of the *contingent* agents  $t$  whose preference depends on the actual world:  $v_t(L, \mathbf{R}) > v_t(L, \mathbf{A})$  and  $v_t(H, \mathbf{A}) > v_t(H, \mathbf{R})$ .

Let  $\alpha_F = \frac{|F|}{|T|}$ ,  $\alpha_U = \frac{|U|}{|T|}$  and  $\alpha_C = \frac{|C|}{|T|}$  be the fractions of the three types of agents. Since the numbers of theory, AI, software, hardware faculty members are known to everyone, we assume that the values of  $\alpha_F$ ,  $\alpha_U$  and  $\alpha_C$  are common knowledge. Admittedly, this assumption may not apply to some specific scenarios. In Appendix D, we discuss the model where agents have only partial information on  $\alpha_F$ ,  $\alpha_U$  and  $\alpha_C$ , and present a strong impossibility result for this model.

The goal is to output the *majority wish*, the alternative that is preferred by at least half of the agents conditioned on the true state. We assume  $T$  is an odd number to avoid ties. Clearly,  $\mathbf{A}$  should be output if  $\alpha_F > \frac{1}{2}$ ,  $\mathbf{R}$  should be output if  $\alpha_U > \frac{1}{2}$ . In the case  $\alpha_F, \alpha_U < \frac{1}{2}$ ,  $\mathbf{A}$  should be output if the actual world is  $H$  and  $\mathbf{R}$  should be output if the actual world is  $L$ .

Our results will sometimes require  $T$ , the number of agents, to be sufficiently large, and it may be helpful to think of  $T \rightarrow \infty$ . However, we will always assume that the parameters of the model:  $B$ ,  $\{P_L, P_H\}$ ,  $\{P_{\ell L}, P_{\ell H}, P_{hL}, P_{hH}\}$ , and  $\{\alpha_F, \alpha_U, \alpha_C\}$ , do not depend on  $T$  in any way.

The traditional social choice setting with agents having predetermined preferences can be viewed as a special case of our model, by setting  $|C| = 0$  (i.e., there is no contingent agent).

For the ease of comprehension, we have used the faculty candidate hiring as a running example for this paper. This can be replaced by any example from most practical scenarios where different types of imperfectly informed voters are voting between two alternatives, including all the examples we mentioned in Section 1.

In our election example in the second paragraph of Section 1, **A** and **R** can represent candidates  $a$  and  $b$  respectively. Correspondingly, for this example,  $L$  and  $H$  can represent “ $a$  is better suited” and “ $b$  is better suited” respectively.  $F$  and  $U$  represent voters aligned to party  $A$  and  $B$  respectively, while  $C$  represents those swing voters whose preferences depend on the signals (in this case, the signals correspond to their private experiences and beliefs, which may be based on information they obtained from their favorite TV programs, newspapers, etc).

## 2.1 Strategy and $\varepsilon$ -strong Bayes Nash equilibrium

A *mechanism* collects a *report* from each agent, and then outputs an alternative which is either **A** or **R**. The mechanism specifies the content of the report by specifying questions for the agents. Examples of those questions include asking each agent for the signal he/she receives, asking each agent to predict the other agents’ reports, etc.

Let  $\mathcal{R}$  be the space of all possible reports, which depends on the design of the mechanism. A *pure strategy* of an agent is given by a function  $\sigma : \mathcal{S} \rightarrow \mathcal{R}$  that maps a signal received by this agent to a report. In a *mixed strategy*,  $\sigma$  can be a random function.

An agent’s strategy is *truthful* if it always specifies the correct answer to each question in the report, to the best of the agent’s knowledge after receiving the signal. For example, if the mechanism asks for the agent’s signal, an agent playing the truthful strategy should report the signal he/she receives; if the mechanism asks the agents to predict the fraction of agents who will receive signal  $m$ , an agent playing the truthful strategy should report his/her posterior belief on this computed by the Bayes rule (Section 2.2 discusses the computation of posterior beliefs).

Given a strategy profile  $\Sigma = (\sigma_1, \dots, \sigma_T)$ , let  $u_t(\Sigma)$  be the expected utility of agent  $t$ , where the expectation is taken over the sampling of agents’ signals. Notice that we use  $u_t$  to denote the *ex-ante* utility (as defined just now) and we have used  $v_t$  to denote the *ex-post* utility (see the fourth paragraph in Section 2). Most parts of this paper will focus on the *ex-ante* utility, especially when we are talking about any equilibrium solution concept.<sup>5</sup>

Since we are in a social choice setting with a potentially large number of agents, a single agent’s behavior may not have much effect. Thus, instead of the typical Bayes Nash Equilibrium, we consider a much stronger goal—the strong Bayes Nash equilibrium.

**Definition 2.1.** A strategy profile  $(\sigma_1, \dots, \sigma_T)$  is an  $\varepsilon$ -strong Bayes Nash equilibrium if there does not exist a subset of agents  $D$  and a strategy profile  $(\sigma'_1, \dots, \sigma'_T)$  such that

1.  $\sigma_t = \sigma'_t$  for each  $t \notin D$ ,
2.  $u_t(\sigma'_1, \dots, \sigma'_T) \geq u_t(\sigma_1, \dots, \sigma_T)$  for each  $t \in D$ , and
3. there exist  $t \in D$  such that  $u_t(\sigma'_1, \dots, \sigma'_T) > u_t(\sigma_1, \dots, \sigma_T) + \varepsilon$ .

When a strategy profile  $(\sigma_1, \dots, \sigma_n)$  is not an  $\varepsilon$ -strong Bayes Nash equilibrium, we will call the subset of the agents  $D$  in Definition 2.1 the *deviating agents* or the *deviating coalition*. By Definition 2.1, every agent of the deviating coalition must be at least as well off, and some must be strictly better off by at least  $\varepsilon$ . Notice that a 0-strong Bayes Nash equilibrium is the conventional strong Bayes Nash equilibrium. The larger  $\varepsilon$  is, the harder it is to find a deviating coalition, and so the larger the set of  $\varepsilon$ -strong Bayes Nash equilibria.

## 2.2 Posterior Update by Bayes Rule

Upon receiving a signal  $S_t \in \{\ell, h\}$ , agent  $t$  updates his/her posterior beliefs (about the probability that (s)he is in world  $L$  or  $H$ , the fraction of agents that will receive signal  $\ell$  or  $h$ , etc.) based on

<sup>5</sup>However, the goal of the mechanism is to output the majority preferred alternative based on agents’ *ex-post* utilities. Our choice of discussing *ex-ante* utilities in solution concept and *ex-post* utilities in the goal is motivated by practical applications. The underlying better alternative is unknown in the election phrase, so *ex-ante* utilities are considered for the solution concept. On the other hand, the ultimate goal is to select the alternative that *turns out to be better* after a certain period of time when the underlying better alternative is revealed (e.g., the performance of a selected president or the effect of a new policy in the next few years after the election), and in this case the *ex-post* utilities are considered.

Bayes rule. Let  $T_{m'm}$  be the probability that an agent who receives signal  $m \in \{\ell, h\}$  believes that another agent will receive signal  $m' \in \{\ell, h\}$ . Straightforward calculations by Bayes rule (available in Appendix B.1) reveals

$$T_{m'm} = \frac{P_L P_{mL}}{P_L P_{mL} + P_H P_{mH}} \cdot P_{m'L} + \frac{P_H P_{mH}}{P_L P_{mL} + P_H P_{mH}} \cdot P_{m'H}. \quad (3)$$

Given a strategy profile  $\Sigma = \{\sigma_1, \dots, \sigma_T\}$  and a mechanism  $\mathcal{M}$ , let  $\lambda_n^{\mathbf{A}, \mathcal{M}}(\Sigma)$  be the probability that alternative  $\mathbf{A}$  is announced as the winner given the actual world is  $n$ , then  $\lambda_n^{\mathbf{R}, \mathcal{M}}(\Sigma) = 1 - \lambda_n^{\mathbf{A}, \mathcal{M}}(\Sigma)$  is the probability that alternative  $\mathbf{R}$  wins given the actual world is  $n$ . We will omit the superscript  $\mathcal{M}$  when it is clear what mechanism we are discussing.

All the agents' *ex-ante* utilities depend exclusively on  $\lambda_L^{\mathbf{A}}(\Sigma), \lambda_H^{\mathbf{A}}(\Sigma)$  (or  $\lambda_L^{\mathbf{R}}(\Sigma), \lambda_H^{\mathbf{R}}(\Sigma)$ ), and each agent  $t$ 's utility is given by

$$u_t(\Sigma) = P_L (\lambda_L^{\mathbf{A}}(\Sigma)v_t(L, \mathbf{A}) + \lambda_L^{\mathbf{R}}(\Sigma)v_t(L, \mathbf{R})) + P_H (\lambda_H^{\mathbf{A}}(\Sigma)v_t(H, \mathbf{A}) + \lambda_H^{\mathbf{R}}(\Sigma)v_t(H, \mathbf{R})), \quad (4)$$

which can also be rewritten as (by noticing  $\lambda_L^{\mathbf{A}}(\Sigma) + \lambda_L^{\mathbf{R}}(\Sigma) = 1$  and  $\lambda_H^{\mathbf{A}}(\Sigma) + \lambda_H^{\mathbf{R}}(\Sigma) = 1$ )

$$u_t(\Sigma) = P_L v_t(L, \mathbf{R}) + P_H v_t(H, \mathbf{R}) + P_L \lambda_L^{\mathbf{A}}(\Sigma)(v_t(L, \mathbf{A}) - v_t(L, \mathbf{R})) + P_H \lambda_H^{\mathbf{A}}(\Sigma)(v_t(H, \mathbf{A}) - v_t(H, \mathbf{R})). \quad (5)$$

We will always use  $\Sigma^* = \{\sigma_1^*, \dots, \sigma_T^*\}$  to denote the truthful strategy profile.

### 3 The Wisdom-of-the-Crowd-Voting Mechanism

We will first review Prelec et al.'s *Surprisingly Popular* algorithm [Prelec et al., 2017], which works under a setting similar to ours but with non-strategic agents. Some part of the intuition behind our mechanism is based on Prelec et al.'s work.

#### 3.1 Prelec et al.'s Surprisingly Popular Algorithm

For the purpose of this paper, we will describe the algorithm with two worlds and two signals. The algorithm asks each agent  $t$  the signal (s)he receives, and his/her belief on the fraction of agents who have received signal  $\ell$  (or signal  $h$ ). In our notation, each agent reports the realization of  $S_t$  and, assuming  $S_t = m \in \{\ell, h\}$ , the value of  $T_{\ell m}$  (or  $T_{hm}$ , which equals to  $1 - T_{\ell m}$ ). Since agents are assumed to be non-strategic, those who receive signal  $\ell$  will report  $(\ell, T_{\ell\ell})$  and those who receive signal  $h$  will report  $(h, T_{\ell h})$ . The algorithm then computes the fraction of agents who report signal  $\ell$ , and the average value of all the reported  $T_{\ell m}$ 's. If the former is greater than the latter,  $\ell$  is considered as being "surprisingly popular" and the algorithm will conclude that  $L$  is the actual world. Otherwise,  $h$  will be considered as being "surprisingly popular" and  $H$  will be concluded as being the actual world.

The correctness of this algorithm is based on the following simple yet important observation in Theorem 3.1. In particular, the average of agents' reported predictions (those  $T_{\ell m}$ 's) will be between  $T_{\ell h}$  and  $T_{\ell\ell}$ . When the number of agents  $T$  is sufficiently large, the actual fraction of agents who receive signal  $\ell$  will be either approximately  $P_{\ell L}$  (if  $L$  is the actual world) or approximately  $P_{\ell H}$  (if  $H$  is the actual world). Theorem 3.1 then implies the correctness of the Surprisingly Popular algorithm.

**Theorem 3.1.**  $P_{\ell H} < T_{\ell h} < T_{\ell\ell} < P_{\ell L}$  and  $P_{hH} > T_{hh} > T_{h\ell} > P_{hL}$

For completeness, the full proof is appears in Appendix B.2, but the intuition is straightforward. The inequality  $P_{\ell H} < P_{\ell L}$  is by the positive correlation (1). The inequality  $T_{\ell h} < T_{\ell\ell}$  is also intuitive: the positive correlation between the signals and worlds implies the positive correlation between two agents' received signals. Finally, (3) implies each of  $T_{\ell h}$  and  $T_{\ell\ell}$  is a weighted average of  $P_{\ell H}$  and  $P_{\ell L}$ , so the value is between  $P_{\ell H}$  and  $P_{\ell L}$ . This concludes the first inequality chain, and the second can be shown similarly.

Throughout this section, we use  $c$  to denote the following constant, which is used in the Chernoff bounds in our proofs.

$$c = \frac{1}{3} \min \{T_{\ell h} - P_{\ell H}, T_{\ell\ell} - T_{\ell h}, P_{\ell L} - T_{\ell\ell}, P_{hH} - T_{hh}, T_{hh} - T_{\ell h}, T_{\ell h} - P_{\ell H}\} \quad (6)$$

### 3.2 The Wisdom-of-the-Crowd-Voting Mechanism

---

#### Mechanism 1 The Wisdom-of-the-Crowd-Voting Mechanism

---

- 1: Each agent  $t$  reports his/her type ( $F$ ,  $U$  or  $C$ ) to the mechanism, and if he/she is of type  $C$ , the signal (s)he receives (either  $\ell$  or  $h$ ), denoted by  $\bar{s}_i \in \{\ell, h\}$ .
  - 2: If agent  $t$  reports type  $F$ , his reported signal will be automatically treated as  $\bar{s}_t = h$ ; if agent  $t$  reports type  $U$ , his reported signal will be automatically treated as  $\bar{s}_t = \ell$ .
  - 3: For each agent  $t$  of type  $C$ , ask him/her to predict the fraction of agents who will report signal  $h$ . Let  $\bar{\delta}_t$  be  $t$ 's prediction. *The prediction  $\bar{\delta}_t$  should be made with the type  $F$  and type  $C$  agents' predictions defined in the previous step being considered, and the mechanism makes this clear to the agents.* For each agent  $t$  of type  $F$ ,  $\bar{\delta}_t$  is set to 0, and for each agent  $t$  of type  $U$ ,  $\bar{\delta}_t$  is set to 1.
  - 4: Compute the *median* of those  $\bar{\delta}_t$ , denoted by  $\bar{\delta}$ .
  - 5: If more than half of the agents report type  $F$ , announce **A** being the winning alternative; if more than half of the agents reports type  $U$ , announce **R** being the winning alternative.
  - 6: If the fraction of the agents reporting  $\bar{s}_t = h$  is more than the median  $\bar{\delta}$ , announce **A** being the winning alternative; otherwise, announce **R** being the winning alternative.
- 

At Step 3 of the mechanism, we only elicit predictions from those contingent agents, and those candidate-friendly (candidate-unfriendly resp.) agents' predictions are treated as 0 (1 resp.). In Appendix C, we discuss an alternative mechanism where we elicit all the predictions and then take the median. The alternative mechanism shares the same theoretical properties, and we discuss the advantages and disadvantages of this alternative.

Mechanism 1 may look too obscure to be implemented in practice. However, very simple and understandable questionnaires implementing the mechanism can be designed. In our running example of faculty candidate hiring, the questionnaire corresponding to our mechanism could look like:

1. Choose one of the followings:
  - (a) I definitely want to accept this candidate, regardless of my colleagues' inputs.
  - (b) I definitely want to reject this candidate, regardless of my colleagues' inputs.
  - (c) After talking to the candidate, I am more inclined to accept him/her than before.
  - (d) After talking to the candidate, I am more inclined to reject him/her than before.
2. If your answer is (c) or (d) in the first question, what percentage of the faculty members do you believe will choose (a) or (c) in the first question?

In the first question above, faculty members of type  $F$  (type  $U$  resp.) will choose (a) ((b) resp.). Faculty members of type  $C$  will choose either (c) or (d) depending on the signals they have received. If a faculty member receives signal  $h$ , (s)he believes world  $H$  is more likely than before. Notice that, it is still possible that (s)he believes the probability of world  $H$  being the actual world is less than 50% and (s)he still prefers rejecting the candidate based on his/her private information (for example, his/her prior belief may be only 10% for world  $H$ , and his/her posterior belief for this increases to 30% upon receiving a signal  $h$ ), so the description that (s)he is "more inclined to accept the candidate than before" accurately implements Mechanism 1. The same holds for those receiving signal  $\ell$ .

### 3.3 Main Theoretical Results for Our Mechanism

We first show that our mechanism indeed achieves (with an exponentially small failure probability) the goal of outputting the alternative favored by the majority, assuming agents are truth-telling.

**Theorem 3.2.** *If all the agents play the truthful strategy  $\Sigma^*$ , then, with probability at least  $1 - 2\exp(-2c^2\alpha_C T)$  (where  $c$  is the constant defined in Eqn. (6)), our mechanism outputs an alternative that is favored by more than half of the agents.*

*Proof.* Suppose all the agents report truthfully. Step 5 of the mechanism guarantees that the majority wish will be announced with probability 1 if either  $\alpha_F > 0.5$  or  $\alpha_U > 0.5$ . It remains to consider the case where we have both  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ .<sup>6</sup> In this case, **A** is favored by the majority if the actual world is  $H$ , and **R** is favored by the majority if the actual world is  $L$ .

---

<sup>6</sup>Recall that we have assume  $T$  is an odd number, so we cannot have  $\alpha_F = 0.5$  or  $\alpha_U = 0.5$ .



If a contingent agent  $t$  receives signal  $h$ , (s)he will believe that a  $T_{hh}$  fraction of agents receive  $h$  (before the treatment at Step 2), and (s)he will report  $\bar{\delta}_t = \alpha_C T_{hh} + \alpha_F$  (after considering the treatment at Step 2). Similarly, a contingent agent receiving signal  $\ell$  will report  $\bar{\delta}_t = \alpha_C T_{h\ell} + \alpha_F$ . We have  $\bar{\delta}_t = 0$  for each candidate-friendly agent and  $\bar{\delta}_t = 1$  for each candidate-unfriendly agent. Since we are considering the case  $\alpha_F, \alpha_U < 0.5$ , the median  $\bar{\delta}$  is in the interval  $[\alpha_C T_{h\ell} + \alpha_F, \alpha_C T_{hh} + \alpha_F]$  (note that  $T_{h\ell} < T_{hh}$  by Theorem 3.1).

Suppose the actual world is  $L$ . The expected fraction of the agents receiving signal  $h$  would be  $P_{hL}$ , and the expected fraction of the agents reporting signal  $h$  (after the treatment at Step 2) would be  $\alpha_C \cdot P_{hL} + \alpha_F$ . By a Chernoff bound, with probability at least  $1 - 2 \exp(-2c^2 \alpha_C T)$ , the fraction of agents reporting signal  $h$  is in the interval  $[\alpha_C \cdot (P_{hL} - c) + \alpha_F, \alpha_C \cdot (P_{hL} + c) + \alpha_F]$ , which is less than  $\alpha_C \cdot T_{h\ell} + \alpha_F \leq \bar{\delta}$  by Theorem 3.1 and (6). Step 6 of the mechanism indicates that  $\mathbf{R}$  will be announced. The analysis for the case where  $H$  is the actual world is similar.  $\square$

Next, we show that the truthful strategy profile is an  $\varepsilon$ -strong Bayes Nash equilibrium of our mechanism for some exponentially small  $\varepsilon$ .

**Theorem 3.3.** *The truthful strategy profile is an  $\varepsilon$ -strong Bayes Nash equilibrium, where  $\varepsilon = (2B^2 + 4B) \exp(-2c^2 \alpha_C T)$ .*

### 3.4 Proof Theorem 3.3 with $T \rightarrow \infty$

We defer the full proof of Theorem 3.3 to Appendix B.3. Here we prove the following limit version, where  $T \rightarrow \infty$ , which illustrates the key features while eliminating the need for both Chernoff bound analyses and for some additional subtle corner cases.

**Theorem 3.4.** *For  $T \rightarrow \infty$ , the truthful strategy profile is a strong Bayes Nash equilibrium.*

We consider three cases: 1)  $\alpha_F > 0.5$ , 2)  $\alpha_U > 0.5$  and 3)  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ .

In the first case, more than half of the agents are candidate-friendly, and  $\mathbf{A}$  will be announced according to Step 5 of the mechanism if these agents report truthfully. The truthful strategy profile forms a strong Bayes Nash equilibrium, as those candidate-friendly agents receive their maximum utilities by truth-telling and the remaining agents are not able to stop the mechanism from outputting  $\mathbf{A}$  regardless of what they report. The analysis for the second case is the analogous to the first case. It remains to consider the third case.

Under the third case,  $\mathbf{A}$  is favored by the majority if the actual world is  $H$ , and  $\mathbf{R}$  is favored by the majority if the actual world is  $L$ . By the same analysis as in the proof of Theorem 3.2, supposing agents report truthfully, we know that  $\mathbf{A}$  will be output with probability 1 (by taking the limit  $T \rightarrow \infty$ ) if the actual world is  $H$ , and  $\mathbf{R}$  will be output if the actual world is  $L$ . Therefore, the contingent agents (type  $C$ ) receive their maximum utilities, and thus have no incentive to deviate from the truthful strategy.

To conclude that truth-telling is a strong Bayes Nash equilibrium, we will show that there is no coalition of deviating agents  $D$  (see the paragraph following Definition 2.1). Let  $\Sigma'$  be the strategy profile after  $D$ 's deviation.

Next, we show that  $D$  cannot contain both a type  $F$  agent and a type  $U$  agent. The following lemma shows that an increase in a type  $F$  agent's (*ex-ante*) utility always results a decrease in a type  $U$  agent's (*ex-ante*) utility, and vice versa. This is obvious if we are dealing with *ex-post* utilities, as a candidate-friendly agent and a candidate-unfriendly agent always want the opposite alternatives. However, this becomes less obvious for *ex-ante* utilities.

**Lemma 3.5.** *Suppose  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ . Let  $\Sigma^*$  be the truthful strategy profile and  $\Sigma'$  be an arbitrary strategy profile. Let  $t_1$  be an arbitrary candidate-friendly agent and  $t_2$  be an arbitrary candidate-unfriendly agent. We have*

- (i) *If  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) > 0$ , then  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) < 0$ .*
- (ii) *If  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) > 0$ , then  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) < 0$ .*

*Proof.* By (5), we have

$$u_t(\Sigma') - u_t(\Sigma^*) = \Gamma_L(v_t(L, \mathbf{A}) - v_t(L, \mathbf{R})) - \Gamma_H(v_t(H, \mathbf{A}) - v_t(H, \mathbf{R})),$$

where  $\Gamma_L = P_L(\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*))$  and  $\Gamma_H = P_H(\lambda_H^{\mathbf{A}}(\Sigma^*) - \lambda_H^{\mathbf{A}}(\Sigma'))$ . Theorem 3.2 implies  $\lambda_L^{\mathbf{A}}(\Sigma^*) = 0$  and  $\lambda_H^{\mathbf{A}}(\Sigma^*) = 1$  (with  $T \rightarrow \infty$ ), which implies  $\Gamma_L, \Gamma_H \geq 0$ . By (2), we have  $v_t(L, \mathbf{A}) - v_t(H, \mathbf{A}) < 0 < v_t(L, \mathbf{R}) - v_t(H, \mathbf{R})$ , which further implies  $v_t(L, \mathbf{A}) - v_t(L, \mathbf{R}) < v_t(H, \mathbf{A}) - v_t(H, \mathbf{R})$ . Since a type  $F$  agent always prefers  $\mathbf{A}$  and a type  $U$  agent always prefers  $\mathbf{R}$ ,

$$\begin{aligned} 0 < v_{t_1}(L, \mathbf{A}) - v_{t_1}(L, \mathbf{R}) &< v_{t_1}(H, \mathbf{A}) - v_{t_1}(H, \mathbf{R}), & \text{and} \\ v_{t_2}(L, \mathbf{A}) - v_{t_2}(L, \mathbf{R}) &< v_{t_2}(H, \mathbf{A}) - v_{t_2}(H, \mathbf{R}) < 0. \end{aligned}$$

Intuitively, this means  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*)$  is more sensitive to  $\Gamma_H$  and  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*)$  is more sensitively to  $\Gamma_L$ . Formally,  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) > 0$  implies  $\Gamma_L > \Gamma_H \frac{v_{t_1}(H, \mathbf{A}) - v_{t_1}(H, \mathbf{R})}{v_{t_1}(L, \mathbf{A}) - v_{t_1}(L, \mathbf{R})} \geq \Gamma_H$ , and  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) \geq 0$  implies  $\Gamma_L \leq \Gamma_H \frac{v_{t_2}(H, \mathbf{A}) - v_{t_2}(H, \mathbf{R})}{v_{t_2}(L, \mathbf{A}) - v_{t_2}(L, \mathbf{R})} \leq \Gamma_H$ . Thus,  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) > 0$  and  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) \geq 0$  cannot be both true. The proof for (ii) is similar.  $\square$

We have seen that the deviating coalition  $D$  cannot contain any contingent agents (since their *ex-ante* utilities have already been maximized). Thus, Lemma 3.5 implies  $D$  can only be comprised of either candidate-friendly agents or candidate-unfriendly agents. Finally, we show that a minority coalition comprised of only candidate-friendly agents or only candidate-unfriendly agents cannot change the outcome by misreporting. We consider candidate-friendly agents without loss of generality.

Suppose  $D$  contains only candidate-friendly agents. Those agents cannot make the mechanism output  $\mathbf{A}$  at Step 5, because fewer than  $1/2$  of the agents report type  $F$  no matter how agents in  $D$  deviate. To maximize the chance for the mechanism to output  $\mathbf{A}$  at Step 6, those candidate-friendly agents would like to maximize the fraction of agents reporting  $h$  and minimize the median of the prediction  $\bar{\delta}$ . However, the mechanism already does this when those candidate-friendly agents play the truthful strategy: their signals are treated as  $h$  at Step 2, and their predictions are treated as 0 at Step 3. The same analysis works when  $D$  contains only candidate-unfriendly agents. This concludes the proof for Theorem 3.4.

In fact, the arguments in the previous paragraph show that: *truth-telling is a dominant strategy for each candidate-friendly agent and each candidate-unfriendly agent.*

## 4 Remarks, Limitations and Future Work

We have assumed agents are Bayesian. Although this assumption may not be completely realistic in practice, we believe the theoretical properties of our mechanism still hold to a certain extent. For example, agents may not exactly predict  $\alpha_C T_{hh} + \alpha_F$  or  $\alpha_C T_{h\ell} + \alpha_F$  in practice, but it is reasonable to assume that their predictions are roughly around these two numbers, or in between. If so, all the theoretical properties will still hold. Prelec et al. [2017] also assume Bayesian agents in the theoretical analysis of the surprisingly popular method, but their empirical experiments with human subjects suggest the method still works in practice.

Our mechanism can be extended to the setting where a fixed fraction  $\tau$  of acceptance votes is required to adopt a policy. For example, in many countries, constitutional amendments requires a  $2/3$  majority to pass. To do this, we only need to change Step 4 of Mechanism 1 such that  $\bar{\delta}$  is the prediction with rank  $\tau T$ , and change Step 5 such that  $\mathbf{A}$  is announced if more than  $\tau$  fraction of agents report type  $F$  and  $\mathbf{R}$  is announced if more than  $1 - \tau$  fraction of agents report type  $U$ .

One limitation is that our mechanism only deals with two alternatives. While this is natural in many scenarios (accept/reject, election with two candidates), extending our results to more than two alternatives is an interesting future direction, but faces a multitude of hurdles: the median technique will not straightforwardly work, the ‘‘surprisingly popular’’ formalism faces impossibility results [Prelec et al., 2017], and Gibbard-Satterthwaite social choice impossibility results apply.

Another interesting future direction is deployment. Our Wisdom-of-the-Crowd-Voting mechanism could be used in any place that currently uses majority voting to better aggregate information. Of course, it would suffer from some of the same drawbacks of majority voting: that the majority can impose their will on the minority. It is not clear if either one of these enjoys fairness properties not included by the other, but that would be another direction of future inquiry. It would be interesting to test this mechanism in the real world, and then test to see if participants are, in aggregate, more pleased with the alternatives selected by this mechanism as compared with those selected by a majority vote.

## Acknowledgement

We gratefully thank Noah Burrell and the anonymous reviewers for the helpful comments and suggestions for improving this paper.

## Funding Transparency Statement

All the funding support related to this paper has been declared at the bottom of Page 1.

## References

- David Austen-Smith and Jeffrey S. Banks. Information aggregation, rationality, and the condorcet jury theorem. *American political science review*, 90(1):34–45, 1996.
- Duncan Black. On the Rationale of Group Decision-making. *Journal of political economy*, 56(1), 1948.
- Kay-Yut Chen, Leslie R Fine, and Bernardo A Huberman. Eliminating public knowledge biases in information-aggregation mechanisms. *Management Science*, 50(7):983–994, 2004.
- Yiling Chen, Chara Podimata, Ariel D Procaccia, and Nisarg Shah. Strategyproof linear regression in high dimensions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 9–26, 2018.
- Anirban Dasgupta and Arpita Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, pages 319–330. International World Wide Web Conferences Steering Committee, 2013.
- N. De Condorcet. *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. L’imprimerie royale, 1785.
- Ofer Dekel, Felix Fischer, and Ariel D Procaccia. Incentive compatible regression learning. *Journal of Computer and System Sciences*, 76(8):759–777, 2010.
- Boi Faltings and Goran Radanovic. Game theory for data science: Eliciting truthful information. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 11(2):1–151, 2017.
- Timothy Feddersen and Wolfgang Pesendorfer. Voting behavior and information aggregation in elections with private information. *Econometrica*, 65(5):1029–1058, 1997. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/2171878>.
- A. Gao, J. R. Wright, and K. Leyton-Brown. Incentivizing Evaluation via Limited Access to Ground Truth: Peer-Prediction Makes Things Worse. *ArXiv e-prints*, June 2016.
- Ralph Hertwig. Tapping into the wisdom of the crowd—with confidence. *Science*, 336(6079):303–304, 2012.
- Hadi Hosseini, Debmalya Mandal, Nisarg Shah, and Kevin Shi. Surprisingly popular voting recovers rankings, surprisingly! In *Proceedings of the International Joint Conference on Artificial Intelligence (forthcoming)*, 2021.
- Harold Hotelling. Stability in Competition. *The Economic Journal*, 39(153), 1929.
- Leslie K John, George Loewenstein, and Drazen Prelec. Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological science*, 23(5):524–532, 2012.
- Yuqing Kong. Dominantly truthful multi-task peer prediction with a constant number of tasks. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2020.
- Yuqing Kong and Grant Schoenebeck. Equilibrium selection in information elicitation without verification via information monotonicity. In *Proceedings of the 9th Innovations in Theoretical Computer Science (ITCS 2018)*, January 2018a.

- Yuqing Kong and Grant Schoenebeck. Eliciting expertise without verification. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, pages 195–212, 2018b.
- Yuqing Kong and Grant Schoenebeck. A Framework For Designing Information Elicitation Mechanisms That Reward Truth-telling. *ACM Transactions on Economics and Computation (TEAC)*, 7(1), 2019.
- Yuqing Kong, Grant Schoenebeck, Biaoshuai Tao, and Fang-Yi Yu. Information elicitation mechanisms for statistical estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34(2), pages 2095–2102, 2020.
- Krishna K. Ladha. The Condorcet Jury Theorem, Free Speech, and Correlated Votes. *American Journal of Political Science*, 36(3), 1992.
- N. Miller. Information, Electorates, and Democracy: some Extensions and Interpretations of the Condorcet Jury Theorem. *Information Pooling and Group Decision Making*, 1986.
- N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, pages 1359–1373, 2005.
- Shmuel Nitzan and Jacob Paroush. Collective decision making and jury theorems. *The Oxford Handbook of Law and Economics*, 1, 2017.
- Javier Perote and Juan Perote-Pena. Strategy-proof estimators for simple regression. *Mathematical Social Sciences*, 47(2):153–176, 2004.
- D. Prelec. A Bayesian Truth Serum for subjective data. *Science*, 306(5695):462–466, 2004.
- Dražen Prelec, H Sebastian Seung, and John McCoy. A solution to the single-question crowd wisdom problem. *Nature*, 541(7638):532–535, 2017.
- Grant Schoenebeck and Fang-Yi Yu. Two strongly truthful mechanisms for three heterogeneous agents answering one question. In *Web and Internet Economics*, pages 119–132, Cham, 2020. Springer International Publishing. ISBN 978-3-030-64946-3.
- Grant Schoenebeck and Fang-Yi Yu. Learning and strongly truthful multi-task setting peer prediction: A variational approach. In *Proceedings of the 12th Innovations in Theoretical Computer Science (ITCS 2021)*, January 2021.
- Victor Shnayder, Arpit Agarwal, Rafael Frongillo, and David C Parkes. Informed truthfulness in multi-task peer prediction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 179–196, 2016.
- Joseph P Simmons, Leif D Nelson, Jeff Galak, and Shane Frederick. Intuitive biases in choice versus estimation: Implications for the wisdom of crowds. *Journal of Consumer Research*, 38(1):1–15, 2011.
- Peyton Young. Condorcet’s Theory of Voting. *American Political Science*, 82(4), 1988.
- Peter Zhang and Yiling Chen. Elicitable and knowledge-free elicitation with peer prediction. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 245–252, 2014.

## A Comparison with Feddersen and Pesendorfer’s Work

Feddersen and Pesendorfer [1997] consider a two-alternative setting similar to our model. As mentioned before, Feddersen and Pesendorfer [1997] consider the standard majority voting where each agent votes for an alternative, while assuming agents play a (Bayes) Nash equilibrium strategy profile. We, on the other hand, design a more sophisticated mechanism to incentivize truth-telling.

Other than this difference, the state space, the signal space and the space of agents’ types in Feddersen and Pesendorfer [1997] are all continuous. For Feddersen and Pesendorfer’s continuous setting, in the Nash equilibrium, agents’ strategies have three types: always vote for one alternative, always vote for the other, and vote the alternative based on the signal. These are similar to our three types: candidate-friendly, contingent, candidate-unfriendly. However, due to continuity, each agent needs to compute his/her type by solving an equation with a complicated Riemann integral (while agents’ know their types directly according to their utility functions in our setting). A phenomenon in their setting due to continuity is that the fraction of contingent voters in the Nash equilibrium approaches zero when the number of the voters goes to infinity.

Although agents can be classified by three types in both settings, we would like to clarify a fundamental difference in the motivation behind this classification. In our setting, each agent’s type reflects his/her preference over the two alternatives. In Feddersen and Pesendorfer’s setting, each agent “chooses” a type in a specific way so that the majority voting scheme outputs the correct alternative with high probability. Therefore, in their setting, agents’ types are chosen for collaboratively aggregating information, and should not be viewed as reflections of their preferences. Although an agent’s preference does affect his/her choice, the purpose for choosing a type is for information aggregation, not for reflecting the preference.

At a high level, our mechanism includes some novel techniques, including the surprisingly popular technique and the median trick, to ensure the output of the correct alternative in the setting with strategic agents. In Feddersen and Pesendorfer’s setting, it may be surprisingly that the simple majority voting scheme is already enough for output the correct alternative. The reason behind this is that certain implicit techniques for guaranteeing the correct output are “embedded” into agents’ strategic behaviors. In other words, the agents are the ones who work out those techniques, not the mechanism. That is why we mentioned in the introduction that the agents in Feddersen and Pesendorfer’s setting need to have much more sophistication compared with the agents in our setting.

Another difference is that they are considering a Nash equilibrium strategy profile, while our mechanism satisfies the much stronger criterion that truth-telling strategies form a *strong* Bayes Nash equilibrium.

## B Omitted Proofs

Equations, theorems, lemmas or propositions are restated for the ease of reading.

### B.1 Proof of Equation (3)

Equation (3) is restated below:

$$T_{m'm} = \frac{P_L P_{mL}}{P_L P_{mL} + P_H P_{mH}} \cdot P_{m'L} + \frac{P_H P_{mH}}{P_L P_{mL} + P_H P_{mH}} \cdot P_{m'H}.$$

*Proof.* Suppose agent  $t$  receives signal  $S_t = m \in \{\ell, h\}$ . (S)he believes that the actual world is  $n \in \{L, H\}$  with probability

$$\Pr(W = n \mid S_t = m) = \frac{\Pr(W = n) \Pr(S_t = m \mid W = n)}{\Pr(S_t = m)} = \frac{P_n P_{mn}}{P_L P_{mL} + P_H P_{mH}}.$$

Now, we can calculate  $T_{m'm}$ :

$$\begin{aligned} T_{m'm} &= P_{m'L} \cdot \Pr(W = L \mid S_t = m) + P_{m'H} \cdot \Pr(W = H \mid S_t = m) \\ &= \frac{P_L P_{mL}}{P_L P_{mL} + P_H P_{mH}} \cdot P_{m'L} + \frac{P_H P_{mH}}{P_L P_{mL} + P_H P_{mH}} \cdot P_{m'H}. \end{aligned}$$

□

## B.2 Proof of Theorem 3.1

**Theorem 3.1.**  $P_{\ell H} < T_{\ell h} < T_{\ell\ell} < P_{\ell L}$  and  $P_{hH} > T_{hh} > T_{h\ell} > P_{hL}$

*Proof.* We will only show the first inequality chain. The second chain follows directly from the first by noticing each term in the second chain is 1 minus a term in the first.

By (3) and  $P_{\ell H} < P_{\ell L}$  in (1), we have

$$T_{\ell\ell} = \frac{P_L P_{\ell L}^2 + P_H P_{\ell H}^2}{P_L P_{\ell L} + P_H P_{\ell H}} < \frac{P_L P_{\ell L}^2 + P_H P_{\ell H} P_{\ell L}}{P_L P_{\ell L} + P_H P_{\ell H}} = P_{\ell L}$$

and

$$T_{\ell h} = \frac{P_L P_{\ell L} P_{hL} + P_H P_{\ell H} P_{hH}}{P_L P_{hL} + P_H P_{hH}} > \frac{P_L P_{\ell H} P_{hL} + P_H P_{\ell H} P_{hH}}{P_L P_{hL} + P_H P_{hH}} = P_{\ell H}.$$

Finally, to show  $T_{\ell\ell} > T_{\ell h}$ , it suffices to show that

$$\pi_1 := \frac{P_L P_{\ell L}}{P_L P_{\ell L} + P_H P_{\ell H}} > \pi_2 := \frac{P_L P_{hL}}{P_L P_{hL} + P_H P_{hH}},$$

since  $T_{\ell\ell} = \pi_1 P_{\ell L} + (1 - \pi_1) P_{\ell H}$ ,  $T_{\ell h} = \pi_2 P_{\ell L} + (1 - \pi_2) P_{\ell H}$  and  $P_{\ell L} > P_{\ell H}$ . Simple calculations show this:

$$\pi_1 > \frac{P_L}{P_L + P_H} > \pi_2,$$

where the first inequality is due to  $P_{\ell L} > P_{\ell H}$  and the second inequality is due to  $P_{hH} > P_{hL}$ .  $\square$

## B.3 Proof of Theorem 3.3

**Theorem 3.3.** *The truthful strategy profile is an  $\varepsilon$ -strong Bayes Nash equilibrium, where  $\varepsilon = (2B^2 + 4B) \exp(-2c^2\alpha_C T)$ .*

### B.3.1 Proof Sketch for Theorem 3.3

We first describe a sketch of the proof. We consider three cases: 1)  $\alpha_F > 0.5$ , 2)  $\alpha_U > 0.5$  and 3)  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ .

**Case 1)**  $\alpha_F > 0.5$ . For the first case, more than half of the agents are candidate-friendly, and **A** will be announced according to Step 5 of the mechanism if these agents report truthfully. The truthful strategy profile forms a (0-)strong Bayes Nash equilibrium, as those candidate-friendly agents receive their maximum utilities by truth-telling and the remaining agents are not able to stop the mechanism from outputting **A** regardless of what they report.

**Case 2)**  $\alpha_U > 0.5$ . The analysis for the second case is the analogous to the first.

**Case 3)**  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ . Under the third case, **A** is favored by the majority if the actual world is **H**, and **R** is favored by the majority if the actual world is **L**. Given a strategy profile  $\Sigma$ , we define

$$I(\Sigma) \equiv P_L \lambda_L^{\mathbf{A}}(\Sigma) + P_H \lambda_H^{\mathbf{R}}(\Sigma)$$

to be the error rate of a strategy: the probability the mechanism selects the alternative that is not the majority wish.

By Theorem 3.2, we know that  $I(\Sigma^*) \leq 2 \exp(-2c^2\alpha_C T)$ . Specifically, supposing agents report truthfully, we know that **A** will be output with probability at least  $1 - 2 \exp(-2c^2\alpha_C T)$  if the actual world is **H** (i.e.,  $\lambda_H^{\mathbf{A}}(\Sigma^*) \geq 1 - 2 \exp(-2c^2\alpha_C T)$ ), and **R** will be output with probability at least  $1 - 2 \exp(-2c^2\alpha_C T)$  if the actual world is **L** (i.e.,  $\lambda_L^{\mathbf{R}}(\Sigma^*) \geq 1 - 2 \exp(-2c^2\alpha_C T)$ ).

To conclude that truth-telling is an  $\varepsilon$ -strong Bayes Nash equilibrium, we will study two cases and show that no coalition of deviating agents **D** exists in either case. Recall that in a deviating coalition all agents must benefit and some agent must benefit by at least  $\varepsilon = (2B^2 + 4B) \exp(-2c^2\alpha_C T)$ . Let  $\Sigma'$  be the strategy profile after **D**'s deviation. We consider two sub-cases:

1.  $I(\Sigma') < (2B + 2) \exp(-2C^2\alpha_C T)$ , and
2.  $I(\Sigma') \geq (2B + 2) \exp(-2C^2\alpha_C T)$ .

In the first case,  $I(\Sigma')$  is small, so the mechanism nearly always chooses the majority wish under  $\Sigma'$ . Therefore, the output of the mechanism does not change with high probability from profile  $\Sigma^*$  to profile  $\Sigma'$ . In particular, we have  $\lambda_H^A(\Sigma') \approx \lambda_H^A(\Sigma^*) \approx 1$  and  $\lambda_L^R(\Sigma') \approx \lambda_L^R(\Sigma^*) \approx 1$ . Thus, no agent can be much better off because all the agents have nearly the same utilities as before. This is formally proved in Claim B.3.

In the second case,  $I(\Sigma')$  is not small, so the mechanism sometimes fails to choose the majority wish under  $\Sigma'$ . Here, the contingent agents do not fare better (Claim B.4) and thus no contingent agent can be in the deviating coalition. The technical key is then to show that  $D$  cannot contain both a candidate-friendly and a candidate-unfriendly agent. Lemma B.1 shows that a significant increase in a candidate-friendly agent's (*ex-ante*) utility always results a decrease in a candidate-unfriendly agent's (*ex-ante*) utility, and vice versa. This is obvious if we are dealing with *ex-post* utilities, as a candidate-friendly agent and a candidate-unfriendly agent always want the opposite alternatives. However, this becomes much less obvious for *ex-ante* utilities.

Thus, any deviating coalition can only be comprised of either candidate-friendly agents or candidate-unfriendly agents. Finally, in Claim B.5, we show that a minority coalition comprised of only candidate-friendly agents cannot change the outcome by misreporting. This concludes the theorem.

### B.3.2 No Win-win Lemma

A key part of the proof is the following lemma which states that it is impossible that predetermined agents of different alternatives both gain from deviating from truthful. As a corollary, any deviating coalition can only contain predetermined agents of one type.

The proof of Lemma B.1 depends on i) truth-telling nearly always selecting the majority wish; and ii) the monotonicity of the *ex-post* utilities  $v_t(\cdot, \mathbf{A})$  and  $v_t(\cdot, \mathbf{R})$  in the first argument (here, we mean  $v_t(L, \mathbf{A}) < v_t(H, \mathbf{A})$  and  $v_t(L, \mathbf{R}) > v_t(H, \mathbf{R})$ ). In particular, Lemma B.1 does not hold if truth-telling is replaced by an arbitrary strategy profile.

**Lemma B.1.** *Suppose  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ . Let  $\Sigma^*$  be the truthful strategy profile and  $\Sigma'$  be an arbitrary strategy profile. Let  $t_1$  be an arbitrary candidate-friendly agent and  $t_2$  be an arbitrary candidate-unfriendly agent. For any  $\Delta \geq 2B \exp(-2c^2\alpha_C T)$ , we have*

- (i) *If  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) > \Delta$ , then  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) < 0$ .*
- (ii) *If  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) > \Delta$ , then  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) < 0$ .*

*Proof.* We will only show (i), as the proof for (ii) is similar.

Since  $v_t(L, \mathbf{A}) < v_t(H, \mathbf{A})$  and  $v_t(L, \mathbf{R}) > v_t(H, \mathbf{R})$  for any agent  $t$  (see (2)), we have  $v_t(L, \mathbf{A}) - v_t(H, \mathbf{A}) < 0 < v_t(L, \mathbf{R}) - v_t(H, \mathbf{R})$ , which further implies

$$v_t(L, \mathbf{A}) - v_t(L, \mathbf{R}) < v_t(H, \mathbf{A}) - v_t(H, \mathbf{R}).$$

Since a candidate-friendly agent always prefers  $\mathbf{A}$  and a candidate-unfriendly agent always prefers  $\mathbf{R}$ , we have

$$\begin{aligned} 0 < v_{t_1}(L, \mathbf{A}) - v_{t_1}(L, \mathbf{R}) < v_{t_1}(H, \mathbf{A}) - v_{t_1}(H, \mathbf{R}), \text{ and} \\ v_{t_2}(L, \mathbf{A}) - v_{t_2}(L, \mathbf{R}) < v_{t_2}(H, \mathbf{A}) - v_{t_2}(H, \mathbf{R}) < 0. \end{aligned} \quad (7)$$

By referring to (5), this intuitively says that  $t_1$ 's utility difference  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*)$  is more sensitive to  $P_H \lambda_H^A$  while  $t_2$ 's utility difference  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*)$  is more sensitive to  $P_L \lambda_L^A$ .

Suppose  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) > \Delta$  as it is assumed in (i). By (5), we have

$$\begin{aligned} P_L(\lambda_L^A(\Sigma') - \lambda_L^A(\Sigma^*))(v_{t_1}(L, \mathbf{A}) - v_{t_1}(L, \mathbf{R})) \\ + P_H(\lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*))(v_{t_1}(H, \mathbf{A}) - v_{t_1}(H, \mathbf{R})) > \Delta. \end{aligned} \quad (8)$$

We consider two cases:  $\lambda_H^A(\Sigma') \geq \lambda_H^A(\Sigma^*)$  and  $\lambda_H^A(\Sigma') < \lambda_H^A(\Sigma^*)$ .

The intuitions for the remaining part of this proof is as follows. In the first case, the probability of outputting  $\mathbf{A}$  (weakly) increases under world  $H$ . Since Theorem 3.2 tells us  $\lambda_H^A(\Sigma^*)$  is already close to

1, the utility gain for  $t_1$  due to the increased probability of outputting  $\mathbf{A}$  under world  $H$  is insignificant, and we must still have  $\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*) > 0$  to ensure  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) > \Delta$ . However, since the probability of outputting  $\mathbf{A}$  increases under both worlds, the utility for  $t_2$  will decrease. In the second case, the probability of outputting  $\mathbf{A}$  decreases under world  $H$ , which reduces the utility for  $t_1$ . To compensate this, the probability of outputting  $\mathbf{A}$  under world  $L$  must increase, in order to ensure  $u_{t_1}(\Sigma') - u_{t_1}(\Sigma^*) > \Delta$ . Moreover, we must have  $P_L(\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*)) > P_H(\lambda_H^{\mathbf{A}}(\Sigma^*) - \lambda_H^{\mathbf{A}}(\Sigma'))$  since the utility difference for agent  $t_1$  is more sensitive to  $P_H \lambda_H^{\mathbf{A}}$ . However, since the utility difference for agent  $t_2$  is more sensitive to  $P_L \lambda_L^{\mathbf{A}}$ , this will reduce the overall utility for  $t_2$ . These are formally proved below.

**Case 1:**  $\lambda_H^{\mathbf{A}}(\Sigma') \geq \lambda_H^{\mathbf{A}}(\Sigma^*)$ . By Theorem 3.2, we have  $\lambda_H^{\mathbf{A}}(\Sigma^*) \geq 1 - 2 \exp(-2c^2 \alpha_C T)$ , which implies  $\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*) \leq 2 \exp(-2c^2 \alpha_C T)$ , which further implies

$$P_H(\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*))(v_{t_1}(H, \mathbf{A}) - v_{t_1}(H, \mathbf{R})) \leq P_H \cdot 2 \exp(-2c^2 \alpha_C T) \cdot B < \Delta.$$

Putting this into (8), we have  $P_L(\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*))(v_{t_1}(L, \mathbf{A}) - v_{t_1}(L, \mathbf{R})) > 0$ , which implies  $\lambda_L^{\mathbf{A}}(\Sigma') > \lambda_L^{\mathbf{A}}(\Sigma^*)$ . We then must have  $u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) < 0$  since we have

$$\begin{aligned} u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) &= P_L(\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*))(v_{t_2}(L, \mathbf{A}) - v_{t_2}(L, \mathbf{R})) \\ &\quad + P_H(\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*))(v_{t_2}(H, \mathbf{A}) - v_{t_2}(H, \mathbf{R})) \end{aligned}$$

by (5),  $\lambda_H^{\mathbf{A}}(\Sigma') \geq \lambda_H^{\mathbf{A}}(\Sigma^*)$  (Case 1 assumption),  $\lambda_L^{\mathbf{A}}(\Sigma') > \lambda_L^{\mathbf{A}}(\Sigma^*)$  (we have just shown),  $v_{t_2}(L, \mathbf{A}) - v_{t_2}(L, \mathbf{R}) < 0$  and  $v_{t_2}(H, \mathbf{A}) - v_{t_2}(H, \mathbf{R}) < 0$  (since  $t_2$  is candidate-unfriendly).

**Case 2:**  $\lambda_H^{\mathbf{A}}(\Sigma') < \lambda_H^{\mathbf{A}}(\Sigma^*)$ . By (8) and  $\Delta > 0$ , we have

$$P_L(\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*))(v_{t_1}(L, \mathbf{A}) - v_{t_1}(L, \mathbf{R})) > P_H(\lambda_H^{\mathbf{A}}(\Sigma^*) - \lambda_H^{\mathbf{A}}(\Sigma'))(v_{t_1}(H, \mathbf{A}) - v_{t_1}(H, \mathbf{R})),$$

which, by the first inequality in (7), further implies

$$P_L(\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*)) > P_H(\lambda_H^{\mathbf{A}}(\Sigma^*) - \lambda_H^{\mathbf{A}}(\Sigma')) > 0.$$

By the second inequality in (7), this implies

$$P_L(\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*))(v_{t_2}(L, \mathbf{A}) - v_{t_2}(L, \mathbf{R})) < P_H(\lambda_H^{\mathbf{A}}(\Sigma^*) - \lambda_H^{\mathbf{A}}(\Sigma'))(v_{t_2}(H, \mathbf{A}) - v_{t_2}(H, \mathbf{R})),$$

which further implies

$$\begin{aligned} u_{t_2}(\Sigma') - u_{t_2}(\Sigma^*) &= P_L(\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*))(v_{t_2}(L, \mathbf{A}) - v_{t_2}(L, \mathbf{R})) \\ &\quad + P_H(\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*))(v_{t_2}(H, \mathbf{A}) - v_{t_2}(H, \mathbf{R})) < 0. \end{aligned}$$

The lemma concludes.  $\square$

**Corollary B.2.** *Suppose  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ . The set of deviating agents  $D$  cannot contain both candidate friendly and candidate unfriendly agents.*

*Proof.* By 3 in Definition 2.1, there must be an agent  $t$  such that  $u_t(\Sigma') - u_t(\Sigma^*) > \varepsilon \geq 2B \exp(-2c^2 \alpha_C T)$ . Assume this agent is candidate-friendly. Then by Lemma B.1, for any candidate-unfriendly agent  $t'$ , we have  $u_{t'}(\Sigma') - u_{t'}(\Sigma^*) < 0$ . Thus no candidate-unfriendly agent can be in the deviating coalition.

An analogous argument works if the benefiting agent is candidate-unfriendly.  $\square$

### B.3.3 Proof of Theorem 3.3

Now we are ready to prove Theorem 3.3.

Suppose this is not the case. There exists a set of deviating agents  $D$  that can deviate from the truthful strategy such that all of them receive utilities that are at least their original utilities and some of them receive utilities that are  $\varepsilon$  strictly higher than their original utilities. Let  $\Sigma'$  be the strategy profile after agents in  $D$  deviate.

We discuss three different cases: 1)  $\alpha_F > 0.5$ , 2)  $\alpha_U > 0.5$  and 3)  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ . Notice that  $n$  being odd implies neither  $\alpha_F$  nor  $\alpha_U$  can be exactly 0.5.



**Case 1:**  $\alpha_F > 0.5$ . If all agents report truthfully,  $\mathbf{A}$  will be announced with probability 1 according to Step 5 of the mechanism. That is,  $\lambda_L^{\mathbf{A}}(\Sigma^*) = \lambda_H^{\mathbf{A}}(\Sigma^*) = 1$ . By 3 of Definition 2.1, there exists  $t \in D$  such that  $u_t(\Sigma') > u_t(\Sigma^*) + \varepsilon$ . Since  $\lambda_L^{\mathbf{A}}(\Sigma')$  and  $\lambda_H^{\mathbf{A}}(\Sigma')$  completely determine each agent's utility, we must have either  $\lambda_L^{\mathbf{A}}(\Sigma') \neq \lambda_L^{\mathbf{A}}(\Sigma^*)$  or  $\lambda_H^{\mathbf{A}}(\Sigma') \neq \lambda_H^{\mathbf{A}}(\Sigma^*)$ . This means either  $\lambda_L^{\mathbf{A}}(\Sigma') < 1$  or  $\lambda_H^{\mathbf{A}}(\Sigma') < 1$ .

Since a candidate-friendly agent's utility is maximized when both  $\lambda_L^{\mathbf{A}}$  and  $\lambda_H^{\mathbf{A}}$  are 1, a candidate-friendly agent's utility will decrease if the strategy profile is switched from  $\Sigma^*$  to  $\Sigma'$ . By 2 of Definition 2.1,  $D$  does not contain any candidate-friendly agent. However, if this is the case, there are still more than half of the agents that will report  $F$  (as  $\alpha_F > 0.5$ ), and  $\mathbf{A}$  will always be announced by Step 5 of the mechanism. We conclude that  $\lambda_L^{\mathbf{A}}(\Sigma') = \lambda_H^{\mathbf{A}}(\Sigma') = 1$ , which contradicts what we have concluded in the previous paragraph.

**Case 2:**  $\alpha_U > 0.5$ . The analysis is similar to the previous case.

**Case 3:**  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ . We consider two sub-cases:  $I(\Sigma') < (2B + 2) \exp(-2c^2\alpha_C T)$  and  $I(\Sigma') \geq (2B + 2) \exp(-2c^2\alpha_C T)$

Firstly, we consider  $I(\Sigma') < (2B + 2) \exp(-2c^2\alpha_C T)$ .

**Claim B.3.** *If  $\alpha_F < 0.5$ ,  $\alpha_U < 0.5$  and  $I(\Sigma') < (2B + 2) \exp(-2c^2\alpha_C T)$ , then  $u_t(\Sigma') - u_t(\Sigma) < \varepsilon$  for every agent  $t$ .*

The ideas behind this proof is that the outcome for  $\Sigma'$  is too close to the outcome of the truthful strategy profile  $\Sigma^*$ , so no agent can get significantly more benefit.

*Proof.* The proof of this claim shows that none of the three types of agents can benefit by  $\varepsilon$  because nothing is substantially different from when agents play truthfully.

By the inequality  $I(\Sigma') = P_L \lambda_L^{\mathbf{A}}(\Sigma') + P_H \lambda_H^{\mathbf{R}}(\Sigma') < (2B + 2) \exp(-2c^2\alpha_C T)$ , we have

$$0 \leq \lambda_H^{\mathbf{R}}(\Sigma') \leq \frac{(2B + 2) \exp(-2c^2\alpha_C T)}{P_H}$$

and

$$0 \leq \lambda_L^{\mathbf{A}}(\Sigma') \leq \frac{(2B + 2) \exp(-2c^2\alpha_C T)}{P_L}.$$

Since  $\lambda_L^{\mathbf{A}}(\Sigma^*) \leq 2 \exp(-2c^2\alpha_C T)$  and  $\lambda_H^{\mathbf{R}}(\Sigma^*) \leq 2 \exp(-2c^2\alpha_C T)$  by Theorem 3.2, we have

$$\lambda_L^{\mathbf{A}}(\Sigma^*) - \lambda_L^{\mathbf{A}}(\Sigma') \leq \lambda_L^{\mathbf{A}}(\Sigma^*) \leq 2 \exp(-2c^2\alpha_C T),$$

$$\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*) \leq \lambda_L^{\mathbf{A}}(\Sigma') \leq \frac{(2B + 2) \exp(-2c^2\alpha_C T)}{P_L},$$

$$\lambda_H^{\mathbf{A}}(\Sigma^*) - \lambda_H^{\mathbf{A}}(\Sigma') \leq 1 - (1 - \lambda_H^{\mathbf{R}}(\Sigma')) \leq \frac{(2B + 2) \exp(-2c^2\alpha_C T)}{P_H},$$

and

$$\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*) \leq 1 - (1 - \lambda_H^{\mathbf{R}}(\Sigma^*)) \leq 2 \exp(-2c^2\alpha_C T).$$

Now, we substitute the four inequalities into the following equation implied by (5):

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

For any candidate-friendly agent, we have  $v_t(L, \mathbf{A}) > v_t(L, \mathbf{R})$  and  $v_t(H, \mathbf{A}) > v_t(H, \mathbf{R})$ , which yields

$$\begin{aligned} &u_t(\Sigma') - u_t(\Sigma^*) \\ &= P_L (\lambda_L^{\mathbf{A}}(\Sigma') - \lambda_L^{\mathbf{A}}(\Sigma^*)) (v_t(L, \mathbf{A}) - v_t(L, \mathbf{R})) + P_H (\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*)) (v_t(H, \mathbf{A}) - v_t(H, \mathbf{R})) \\ &\leq P_L \cdot \frac{(2B + 2) \exp(-2c^2\alpha_C T)}{P_L} \cdot B + P_H \cdot 2 \exp(-2c^2\alpha_C T) \cdot B \\ &< (2B^2 + 4B) \exp(-2c^2\alpha_C T) = \varepsilon. \end{aligned}$$

For any contingent agent, we have  $v_t(L, \mathbf{R}) > v_t(L, \mathbf{A})$  and  $v_t(H, \mathbf{A}) > v_t(H, \mathbf{R})$ , which yields

$$\begin{aligned} & u_t(\Sigma') - u_t(\Sigma^*) \\ &= P_L(\lambda_L^{\mathbf{A}}(\Sigma^*) - \lambda_L^{\mathbf{A}}(\Sigma'))(v_t(L, \mathbf{R}) - v_t(L, \mathbf{A})) + P_H(\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*))(v_t(H, \mathbf{A}) - v_t(H, \mathbf{R})) \\ &\leq P_L \cdot 2 \exp(-2c^2\alpha_C T) \cdot B + P_H \cdot 2 \exp(-2c^2\alpha_C T) \cdot B \\ &= 2B \exp(-2c^2\alpha_C T) < \varepsilon. \end{aligned}$$

For any candidate-unfriendly agent, we have  $v_t(L, \mathbf{R}) > v_t(L, \mathbf{A})$  and  $v_t(H, \mathbf{R}) > v_t(H, \mathbf{A})$ , which yields

$$\begin{aligned} & u_t(\Sigma') - u_t(\Sigma^*) \\ &= P_L(\lambda_L^{\mathbf{A}}(\Sigma^*) - \lambda_L^{\mathbf{A}}(\Sigma'))(v_t(L, \mathbf{R}) - v_t(L, \mathbf{A})) + P_H(\lambda_H^{\mathbf{A}}(\Sigma^*) - \lambda_H^{\mathbf{A}}(\Sigma'))(v_t(H, \mathbf{R}) - v_t(H, \mathbf{A})) \\ &\leq P_L \cdot 2 \exp(-2c^2\alpha_C T) \cdot B + P_H \cdot \frac{(2B + 2) \exp(-2c^2\alpha_C T)}{P_H} \cdot B \\ &< (2B^2 + 4B) \exp(-2c^2\alpha_C T) = \varepsilon. \end{aligned}$$

We conclude none of the agents has a utility gain of at least  $\varepsilon$ , which contradict 3 of Definition 2.1.  $\square$

Claim B.3 implies that, in the first case  $I(\Sigma') < (2B + 2) \exp(-2c^2\alpha_C T)$ , there does not exist a deviating coalition  $D$  where an agent in  $D$  can receive an utility gain of at least  $\varepsilon$ , which contradicts to our assumption about  $D$  at the beginning.

Next, we consider the second case  $I(\Sigma') \geq (2B + 2) \exp(-2c^2\alpha_C T)$ . This case is more complicated. We first show that no contingent agent in  $\Sigma'$  can do as well as in the truthful profile  $\Sigma^*$ .

**Claim B.4.** *If  $\alpha_F < 0.5$ ,  $\alpha_U < 0.5$  and  $I(\Sigma') \geq (2B + 2) \exp(-2c^2\alpha_C T)$ , then  $u_t(\Sigma') - u_t(\Sigma^*) < 0$  for every contingent agent  $t$ .*

The ideas behind this proof is that those contingent agents already receive almost optimal utilities in  $\Sigma^*$ ; therefore, if the error rate of the strategy  $\Sigma'$  is high enough, the utilities of the contingent agents will decrease.

*Proof.* By Theorem 3.2, we have

$$\begin{aligned} P_L \lambda_L^{\mathbf{R}}(\Sigma^*) + P_H \lambda_H^{\mathbf{A}}(\Sigma^*) &\geq P_L(1 - 2 \exp(-2c^2\alpha_C T)) + P_H(1 - 2 \exp(-2c^2\alpha_C T)) \\ &= 1 - 2 \exp(-2c^2\alpha_C T). \end{aligned} \quad (9)$$

Therefore

$$\begin{aligned} P_L \lambda_L^{\mathbf{R}}(\Sigma^*) + P_H \lambda_H^{\mathbf{A}}(\Sigma^*) &\geq 1 - 2 \exp(-2c^2\alpha_C T) \\ &= 2B \exp(-2c^2\alpha_C T) + 1 - (2B + 2) \exp(-2c^2\alpha_C T) \\ &\geq 2B \exp(-2c^2\alpha_C T) + 1 - I(\Sigma') \\ &= P_L \lambda_L^{\mathbf{R}}(\Sigma') + P_H \lambda_H^{\mathbf{A}}(\Sigma') + 2B \exp(-2c^2\alpha_C T) \end{aligned} \quad (10)$$

By rewriting (4) as (noticing  $\lambda_L^{\mathbf{A}}(\Sigma) + \lambda_H^{\mathbf{A}}(\Sigma) = 1$ )

$$\begin{aligned} u_t(\Sigma) &= P_L v_t(L, \mathbf{A}) + P_H v_t(H, \mathbf{R}) + P_L \lambda_L^{\mathbf{R}}(\Sigma)(v_t(L, \mathbf{R}) - v_t(L, \mathbf{A})) \\ &\quad + P_H \lambda_H^{\mathbf{A}}(\Sigma)(v_t(H, \mathbf{A}) - v_t(H, \mathbf{R})), \end{aligned}$$

we have

$$\begin{aligned} u_t(\Sigma') - u_t(\Sigma^*) &= P_L(v_t(L, \mathbf{R}) - v_t(L, \mathbf{A}))(\lambda_L^{\mathbf{R}}(\Sigma') - \lambda_L^{\mathbf{R}}(\Sigma^*)) \\ &\quad + P_H(v_t(H, \mathbf{A}) - v_t(H, \mathbf{R}))(\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*)) \end{aligned} \quad (11)$$

We will show  $u_t(\Sigma') - u_t(\Sigma^*) < 0$  for an arbitrary contingent agent  $t$ . Recall that  $v_t(L, \mathbf{R}) - v_t(L, \mathbf{A}) > 0$  and  $v_t(H, \mathbf{A}) - v_t(H, \mathbf{R}) > 0$ . We consider three cases:

- If  $\lambda_H^{\mathbf{A}}(\Sigma') \leq \lambda_H^{\mathbf{A}}(\Sigma^*)$  and  $\lambda_L^{\mathbf{R}}(\Sigma') \leq \lambda_L^{\mathbf{R}}(\Sigma^*)$ , then one of these two inequalities must be strict by (10). Equation (11) then implies  $u_t(\Sigma') - u_t(\Sigma^*) < 0$ .

- If  $\lambda_H^{\mathbf{A}}(\Sigma') > \lambda_H^{\mathbf{A}}(\Sigma^*)$ , then we have  $P_L \lambda_L^{\mathbf{R}}(\Sigma^*) - P_L \lambda_L^{\mathbf{R}}(\Sigma') > 2B \exp(-2c^2 \alpha_C T)$  by (10). Since  $\lambda_H^{\mathbf{A}}(\Sigma') \leq 1$  and  $\lambda_H^{\mathbf{A}}(\Sigma^*) \geq 1 - 2 \exp(-2c^2 \alpha_C T)$ , we have  $\lambda_H^{\mathbf{A}}(\Sigma') - \lambda_H^{\mathbf{A}}(\Sigma^*) \leq 2 \exp(-2c^2 \alpha_C T)$ . We also have  $v_t(H, \mathbf{A}) - v_t(H, \mathbf{R}) \leq B$  and  $v_t(L, \mathbf{R}) - v_t(L, \mathbf{A}) \geq 1$  (recall that  $v_t(L, \mathbf{A})$ ,  $v_t(L, \mathbf{R})$ ,  $v_t(H, \mathbf{R})$  and  $v_t(H, \mathbf{A})$  are integers bounded by  $B$ ). Putting those into (11), we have

$$u_t(\Sigma') - u_t(\Sigma^*) < 1 \cdot (-2B \exp(-2c^2 \alpha_C T)) + P_H \cdot B \cdot 2 \exp(-2c^2 \alpha_C T) < 0.$$

- If  $\lambda_L^{\mathbf{R}}(\Sigma') > \lambda_L^{\mathbf{R}}(\Sigma^*)$ , then we have  $P_H \lambda_H^{\mathbf{A}}(\Sigma^*) - P_H \lambda_H^{\mathbf{A}}(\Sigma') > 2B \exp(-2c^2 \alpha_C T)$  by (10). Similar to the second case, we have  $\lambda_L^{\mathbf{R}}(\Sigma') - \lambda_L^{\mathbf{R}}(\Sigma^*) \leq 2 \exp(-2c^2 \alpha_C T)$ ,  $v_t(H, \mathbf{A}) - v_t(H, \mathbf{R}) \geq 1$  and  $v_t(L, \mathbf{R}) - v_t(L, \mathbf{A}) \leq B$ . Putting those into (11), we have

$$u_t(\Sigma') - u_t(\Sigma^*) < P_L \cdot B \cdot 2 \exp(-2c^2 \alpha_C T) + 1 \cdot (-2B \exp(-2c^2 \alpha_C T)) < 0.$$

Putting these three cases together, we have  $u_t(\Sigma') - u_t(\Sigma^*) < 0$  for an arbitrary contingent agent  $t$ .  $\square$

Therefore,  $D$  cannot contain any contingent agents by 2 of Definition 2.1. Corollary B.2 says that  $D$  cannot simultaneously contain an candidate-friendly agent and a candidate-unfriendly agent. Thus any  $D$  must contain either only candidate-friendly agents or only candidate-unfriendly agents.

The following claim states that neither type of predetermined agents alone are powerful enough to change the outcome to their favor. This concludes the proof as we have shown there is no deviating coalition.

**Claim B.5.** *Suppose  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ . If  $D$  contains only candidate-friendly agents, then  $\lambda_L^{\mathbf{A}}(\Sigma') \leq \lambda_L^{\mathbf{A}}(\Sigma^*)$  and  $\lambda_H^{\mathbf{A}}(\Sigma') \leq \lambda_H^{\mathbf{A}}(\Sigma^*)$ . If  $D$  contains only candidate-unfriendly agents, then  $\lambda_L^{\mathbf{R}}(\Sigma') \leq \lambda_L^{\mathbf{R}}(\Sigma^*)$  and  $\lambda_H^{\mathbf{R}}(\Sigma') \leq \lambda_H^{\mathbf{R}}(\Sigma^*)$ .*

*Proof.* Consider candidate-friendly agents without loss of generality. Since contingent agents and candidate-unfriendly agents, which constitute more than half of the population, are truth-telling, those candidate-friendly agents cannot make the mechanism announce  $\mathbf{A}$  at Step 5, since they cannot make more than half of agents report type  $F$ . To maximize the probability that the mechanism announce  $\mathbf{A}$  at Step 6, those candidate-friendly agents would like to maximize the fraction of agents reporting signal  $h$  and minimize the median  $\bar{\delta}$ . However, the mechanism's conversion of signals (Step 2) and predictions (Step 3) already does these for candidate-friendly agents.  $\square$

Thus, we have proved that, for the second case  $I(\Sigma') \geq (2B + 2) \exp(-2c^2 \alpha_C T)$ , no such deviating set  $D$  exists, which contradicts to our assumption for the existence of  $D$  at the beginning.

## C An Alternative Mechanism

As we have remarked right below Mechanism 1, we present an alternative mechanism that achieves the same theoretical properties. The mechanism is shown in Mechanism 2.

---

**Mechanism 2** The Wisdom-of-the-Crowd-Voting Mechanism (an alternative)

- 1: Each agent  $t$  reports to the mechanism the signal (s)he receives (either  $\ell$  or  $h$ ), denoted by  $\bar{s}_t \in \{\ell, h\}$ , his/her type ( $F$ ,  $U$  or  $C$ ), his/her posterior belief of the fraction of agents who will report signal  $h$ , denoted by  $\bar{\delta}_t$ .
  - 2: If agent  $t$  reports type  $F$ , his reported signal will be automatically treated as  $\bar{s}_t = h$ ; if agent  $t$  reports type  $U$ , his reported signal will be automatically treated as  $\bar{s}_t = \ell$ . *The prediction  $\bar{\delta}_t$  in the previous step should be made with this treatment being considered, and the mechanism makes this clear to the agents.*
  - 3: Compute the *median* of the reported  $\bar{\delta}_t$ , denoted by  $\bar{\delta}$ .
  - 4: If more than half of the agents report type  $F$ , announce  $\mathbf{A}$  being the winning alternative; if more than half of the agents reports type  $U$ , announce  $\mathbf{R}$  being the winning alternative.
  - 5: If the number of agents reporting  $\bar{s}_t = h$  is more than the median  $\bar{\delta}$ , announce  $\mathbf{A}$  being the winning alternative; otherwise, announce  $\mathbf{R}$  being the winning alternative.
- 

Correspondingly, the questionnaire becomes the followings.

1. Choose one of the followings:
  - (a) I definitely want to accept this candidate.
  - (b) I definitely want to reject this candidate.
  - (c) After talking to the candidate, I am more inclined to accept him/her than before.
  - (d) After talking to the candidate, I am more inclined to reject him/her than before.
2. What percentage of the faculty members do you believe will choose (a) or (c) in the first question?

Theorem 3.2 still holds for Mechanism 2. If  $\alpha_F > 0.5$  or  $\alpha_U > 0.5$ , the mechanism outputs the majority wish (**A** or **R** respectively) at Step 4 with probability 1 as before. If  $\alpha_F, \alpha_U < 0.5$ , we still have  $\bar{\delta} \in [\alpha_C T_{h\ell} + \alpha_F, \alpha_C T_{hh} + \alpha_F]$ . This is actually easier to see: agents' predictions are now either  $\alpha_C T_{h\ell} + \alpha_F$  or  $\alpha_C T_{hh} + \alpha_F$ . The remaining part of the proof is the same as before.

Theorem 3.3 still holds for Mechanism 2. In fact, all parts of the proof are the same as before, except for Claim B.5 where we can only prove a weaker statement, which is, nevertheless, sufficient to show Theorem 3.3.

**Claim C.1.** *Suppose  $\alpha_F < 0.5$  and  $\alpha_U < 0.5$ . If  $D$  contains only candidate-friendly agents, then  $\lambda_L^A(\Sigma') \leq \lambda_L^A(\Sigma^*) + 2 \exp(-2c^2 \alpha_C T)$  and  $\lambda_H^A(\Sigma') \leq \lambda_H^A(\Sigma^*) + 2 \exp(-2c^2 \alpha_C T)$ . If  $D$  contains only candidate-unfriendly agents, then  $\lambda_L^R(\Sigma') \leq \lambda_L^R(\Sigma^*) + 2 \exp(-2c^2 \alpha_C T)$  and  $\lambda_H^R(\Sigma') \leq \lambda_H^R(\Sigma^*) + 2 \exp(-2c^2 \alpha_C T)$ .*

*Proof.* We focus on the case that  $D$  contains only candidate-friendly agents. The candidate-unfriendly case is analogous.

First of all, since there are strictly less than half of the agents reporting type  $F$  (those type  $U$  and type  $C$  agents, which contribute more than half of the population, report their types truthfully), those candidate-friendly agents in  $D$  cannot make the mechanism output **A** at Step 5 of the mechanism. Therefore, they can only attempt to make the mechanism output **A** at Step 6 with a higher probability.

Suppose the actual world is  $H$ . We need to prove  $\lambda_H^A(\Sigma') \leq \lambda_H^A(\Sigma^*) + 2 \exp(-2c^2 \alpha_C T)$ . This is trivial: Theorem 3.2 implies  $\lambda_H^A(\Sigma^*) \geq 1 - 2 \exp(-2c^2 \alpha_C T)$ , so the right-hand side of the inequality is at least 1, making the inequality always hold.

Suppose the actual world is  $L$ . We need to prove  $\lambda_L^A(\Sigma') \leq \lambda_L^A(\Sigma^*) + 2 \exp(-2c^2 \alpha_C T)$ . It suffices to show  $\lambda_L^A(\Sigma') \leq 2 \exp(-2c^2 \alpha_C T)$ , which is equivalent to

$$\lambda_L^R(\Sigma') \geq 1 - 2 \exp(-2c^2 \alpha_C T). \quad (12)$$

Supposing  $\Sigma'$  is played, we will prove the following two observations:

1. With probability at least  $1 - 2 \exp(-2c^2 \alpha_C T)$ , the fraction of agents reporting signal  $h$  is at most  $\alpha_C (P_{hL} + c) + \alpha_F$ ;
2. The median of the prediction  $\bar{\delta}$  falls into the interval  $[\alpha_C T_{h\ell} + \alpha_F, \alpha_C T_{hh} + \alpha_F]$  (with probability 1).

To show the first observation, all the candidate-unfriendly agents will report signal  $\ell$  (after the conversion in Step 2). For the contingent agents, each of them receives signal  $h$  with probability  $P_{hL}$ . By a Chernoff bound, with probability at least  $1 - 2 \exp(-2c^2 \alpha_C T)$ , the fraction of the contingent agents receiving  $h$  is at most  $P_{hL} + c$ . Even if all the candidate-friendly agents report  $h$ , the overall fraction of agents reporting  $h$  is at most  $\alpha_C (P_{hL} + c) + \alpha_F$  with probability at least  $1 - 2 \exp(-2c^2 \alpha_C T)$ .

To show the second observation, the prediction  $\bar{\delta}_t$  that a truthful agent will report is either  $\alpha_C T_{h\ell} + \alpha_F$  (if (s)he receive signal  $\ell$ ) or  $\alpha_C T_{hh} + \alpha_F$  (if (s)he receive signal  $h$ ). Since there are more than half of truth-telling agents, the median  $\bar{\delta}$  is always within the interval  $[\alpha_C T_{h\ell} + \alpha_F, \alpha_C T_{hh} + \alpha_F]$ .

Finally, by noticing  $\alpha_C (P_{hL} + c) + \alpha_F < \alpha_C T_{h\ell} + \alpha_F$  (implied by Theorem 3.1 and (6)), the two observations imply that the fraction of agents reporting signal  $h$  is strictly less than  $\bar{\delta}$  with probability at least  $1 - 2 \exp(-2c^2 \alpha_C T)$ , which implies (12) by our design in Step 6 of the mechanism.  $\square$

To conclude Theorem 3.3, for each agent  $t$  in  $D$  that contains only candidate-friendly agents, we have

$$\begin{aligned}
& u_t(\Sigma') - u_t(\Sigma^*) \\
&= P_L(\lambda_L^A(\Sigma') - \lambda_L^A(\Sigma^*))(v_t(L, \mathbf{A}) - v_t(L, \mathbf{R})) + P_H(\lambda_H^A(\Sigma') - \lambda_H^A(\Sigma^*))(v_t(H, \mathbf{A}) - v_t(H, \mathbf{R})) \\
&\hspace{15em} \text{(by (5))} \\
&\leq P_L \cdot 2 \exp(-2c^2\alpha_C T) \cdot B + P_H \cdot 2 \exp(-2c^2\alpha_C T) \cdot B \hspace{5em} \text{(by Claim C.1)} \\
&= 2B \exp(-2c^2\alpha_C T) < \varepsilon.
\end{aligned}$$

Thus, no agent in  $D$  satisfies 3 in Definition 2.1. A similar analysis holds for the case where  $D$  contains only candidate-unfriendly agents.

### C.1 Comparison of the Two Mechanisms

Both Mechanism 1 and Mechanism 2 achieve the same set of theoretical properties.

The advantage for Mechanism 1 is that it is “slightly more truthful” in the sense that Claim B.5 is stronger than Claim C.1. In fact, under Mechanism 1, we have seen that truth-telling is a dominant strategy for both candidate-friendly and candidate-unfriendly agents, while this nice property is lost in Mechanism 2. Under Mechanism 2, the dominant strategy for a candidate-friendly agent (candidate-unfriendly agent resp.) is to report prediction 0 (1 resp.), which is no longer a truthful strategy. Nevertheless, we have seen that the truthful strategy is good enough so that a deviation to the dominant strategy does not provide a utility gain of at least  $\varepsilon$ .

Mechanism 2 wins by a little bit for its simplicity and symmetry. It is easier to explain Mechanism 2 to the users in practice. Notice that Mechanism 1 essentially converts the prediction reported from each candidate-friendly agent (candidate-unfriendly agent resp.) to 0 (1 resp.). Converting the predictions may seem to be less natural than converting the signals for users. Especially, for those users who are not familiar to the idea of “surprisingly popular”, they may not be able to realize that converting their predictions to the opposite extreme is helpful for them, and they may be more skeptical of Mechanism 1 due to this. In addition, Mechanism 2 treats the reported predictions symmetrically, which may be more acceptable to the users in practical implementations.

Another related question is how exactly to phrase the ballot in practice. Here, we mimic the questions of Prelec et al. [2017] and ask for a forecast. However, it may be preferable in practice to ask, as in Mechanism 3, for a fractional threshold of (a) and (c) responses above which the agents would prefer to accept. While the outcomes would be mathematically equivalent, one or the other or a third alternative might work better in practice. This is, however, beyond the scope of this paper.

## D Unknown/Partially Known Distribution of Agent Types

Before generalizing our result to the setting with general non-binary signals and states, we take a detour and discuss the necessity of the assumption that the distribution of agent types is common knowledge. As we mentioned in Section 2, we assume the distribution of agent types,  $\alpha_F, \alpha_U, \alpha_C$ , is a common knowledge among the agents. This assumption is natural by its own in many scenarios including our candidate hiring example (if a theory candidate is applying at a computer science department, those theory faculty members are more inclined to accept the candidate than the faculty members in AI, software, hardware; moreover, the numbers of the theory, AI, software, hardware faculty members are public information). In this section, we will see that this assumption is also necessary for the existence of a mechanism that satisfies Theorem 3.2 and Theorem 3.3.

Before describing our impossibility result, we first formally define the model with unknown agent types. Let  $\Delta_3 = \{(x_1, x_2, x_3) \mid \forall i : x_i \in [0, 1]; x_1 + x_2 + x_3 = 1\}$ . The distribution of types,  $(\alpha_F, \alpha_U, \alpha_C)$ , is then an element of  $\Delta_3$ . To model an unknown/partially known distribution of agent types, let  $\mathcal{D}_{\Delta_3}$  be a distribution over  $\Delta_3$  where each agent believes the distribution of the agent types,  $(\alpha_F, \alpha_U, \alpha_C)$ , is drawn from  $\mathcal{D}_{\Delta_3}$ .

Note that while the fraction of types is not known, the prior over the worlds,  $P_L, P_H$ , and the signal structures conditioned on types,  $P_{\ell L}, P_{\ell H}, P_{hL}, P_{hH}$  are still common knowledge.

Next, we describe a natural property that is shared by most social choice mechanism, including the one in this paper.

**Definition D.1.** A mechanism is *anonymous* if it always outputs the same alternative for any two collections of reports  $\mathbf{r}^{(1)} = (r_1^{(1)}, \dots, r_T^{(1)}) \in \mathcal{R}^T$ ,  $\mathbf{r}^{(2)} = (r_1^{(2)}, \dots, r_T^{(2)}) \in \mathcal{R}^T$  such that  $\mathbf{r}^{(1)}$  is a permutation of  $\mathbf{r}^{(2)}$ .

In other words, an anonymous mechanism cannot decide the output alternative based on agents' identities.

We have the following strong impossibility result.

**Theorem D.2.** *Under the setting with an unknown distribution of agent types, there exists a constant  $\tau > 0$  such that no anonymous mechanism always outputs the alternative favored by more than half of the agents with probability more than  $1 - \tau$  in any  $\tau$ -strong symmetric Bayes Nash equilibrium.*

Since the truthful strategy profile is symmetric, we have the following corollary about the impossibility of a truthful mechanism.

**Corollary D.3.** *Under the setting with unknown distribution of agent types defined above, there exists a constant  $\tau > 0$  such that no anonymous mechanism satisfies both of the followings:*

- *the mechanism outputs the alternative favored by more than half of the agents with probability more than  $1 - \tau$ ;*
- *under the mechanism, the truthful strategy profile is a  $\tau$ -strong Bayes Nash equilibrium.*

## D.1 Proof of Theorem D.2

Consider an anonymous mechanism and an arbitrary symmetric strategy profile  $\Sigma$ . Let  $\beta_{F,\ell}$  be the fraction of candidate-friendly agents that receive signal  $\ell$ . Let  $\beta_{F,h}, \beta_{U,\ell}, \beta_{U,h}, \beta_{C,\ell}$  and  $\beta_{C,h}$  have similar meanings. The mechanism can only see how many different reports there are, and how many agents report each of them; in particular, the mechanism cannot see who reports which. The following proposition follows immediately from the above remarks.

**Proposition D.4.** *Fix a symmetric strategy profile  $\Sigma$ . If a mechanism is anonymous, then the values  $\{\alpha_F, \alpha_U, \alpha_C\} \cup \{\beta_{F,\ell}, \beta_{F,h}, \beta_{U,\ell}, \beta_{U,h}, \beta_{C,\ell}, \beta_{C,h}\}$  completely determine the output of the mechanism.*

We also need the following technical lemma.

**Lemma D.5.** *The total variation distance between the two binomial distributions  $\text{Bin}(T, 1/6)$  and  $\text{Bin}(T/3, 1/2)$  is less than 0.123 for sufficiently large  $T$ .*

*Proof.* By Central Limit Theorem, the total variation distance between  $\text{Bin}(T, 1/6)$  and  $\text{Bin}(T/3, 1/2)$  is at most the total variation distance between  $\mathcal{N}(T/6, 5T/36)$  and  $\mathcal{N}(T/6, T/12)$  plus  $o(1)$ , which, by shifting the mean of the Gaussian distribution, is the total variation distance between  $\mathcal{N}(0, 5T/36)$  and  $\mathcal{N}(0, T/12)$  plus  $o(1)$ .

Let  $f(x)$  and  $g(x)$  be the probability density function for  $\mathcal{N}(0, 5T/36)$  and  $\mathcal{N}(0, T/12)$  respectively. To calculate the total variation distance, firstly, straightforward calculations reveal that  $f(x) < g(x)$  on  $(-\sqrt{\frac{5}{24}} \ln \frac{5}{3} T, \sqrt{\frac{5}{24}} \ln \frac{5}{3} T)$  and  $f(x) > g(x)$  on  $(-\infty, -\sqrt{\frac{5}{24}} \ln \frac{5}{3} T) \cup (\sqrt{\frac{5}{24}} \ln \frac{5}{3} T, \infty)$ . Therefore, the total variation distance between  $\mathcal{N}(0, 5T/36)$  and  $\mathcal{N}(0, T/12)$  is

$$\begin{aligned} \int_{-\sqrt{\frac{5}{24}} \ln \frac{5}{3} T}^{\sqrt{\frac{5}{24}} \ln \frac{5}{3} T} g(x) - f(x) dx &= \int_{-\sqrt{\frac{5}{24}} \ln \frac{5}{3} T}^{\sqrt{\frac{5}{24}} \ln \frac{5}{3} T} \frac{1}{\sqrt{2\pi \frac{T}{12}}} e^{-\frac{1}{2} \frac{x^2}{T/12}} - \frac{1}{\sqrt{2\pi \frac{5T}{36}}} e^{-\frac{1}{2} \frac{x^2}{5T/36}} dx \\ &= \int_{-\sqrt{\frac{5}{24}} \ln \frac{5}{3} T}^{\sqrt{\frac{5}{24}} \ln \frac{5}{3} T} \frac{1}{\sqrt{2\pi \frac{1}{12}}} e^{-\frac{1}{2} \frac{y^2}{1/12}} - \frac{1}{\sqrt{2\pi \frac{5}{36}}} e^{-\frac{1}{2} \frac{y^2}{5/36}} dy \\ &\qquad\qquad\qquad (\text{where } y = x/\sqrt{T}) \\ &< 0.12295. \qquad\qquad\qquad (\text{Calculated by computer}) \end{aligned}$$

Thus, the total variation distance between  $\text{Bin}(T, 1/6)$  and  $\mathcal{N}(T/6, 5T/36)$  is at most  $0.12295 + o(1)$ , which implies the lemma.  $\square$

Now we are ready to present the proof of Theorem D.2 (restated below).

**Theorem D.2.** *Under the setting with an unknown distribution of agent types, there exists a constant  $\tau > 0$  such that no anonymous mechanism always outputs the alternative favored by more than half of the agents with probability more than  $1 - \tau$  in any  $\tau$ -strong symmetric Bayes Nash equilibrium.*

*Proof.* We consider the following instance.

The prior distribution of the two worlds (world  $L$  and world  $H$ ) is given by  $P_L = 0.98$  and  $P_H = 0.02$ . The probability distribution of the two signals under each of the two worlds is given by  $P_{\ell L} = 1$ ,  $P_{\ell H} = 0$ ,  $P_{hL} = 5/6$ ,  $P_{hH} = 1/6$ . For each candidate-friendly agent  $t$ , we have  $v_t(H, \mathbf{A}) = 3$ ,  $v_t(H, \mathbf{R}) = 0$ ,  $v_t(L, \mathbf{A}) = 2$  and  $v_t(L, \mathbf{R}) = 1$ . For each contingent agent  $t$  we have  $v_t(H, \mathbf{A}) = 3$ ,  $v_t(H, \mathbf{R}) = 0$ ,  $v_t(L, \mathbf{A}) = 1$  and  $v_t(L, \mathbf{R}) = 2$ . For each candidate-unfriendly agent  $t$ , we have  $v_t(H, \mathbf{A}) = 1$ ,  $v_t(H, \mathbf{R}) = 2$ ,  $v_t(L, \mathbf{A}) = 0$  and  $v_t(L, \mathbf{R}) = 3$ . Lastly,  $\mathcal{D}_{\Delta_3}$  is defined as follows: with probability  $1/2$  we are in setting  $X$  and the fractions of agents with types  $F, C, U$  are  $\alpha_F^{(1)} = 1/3$ ,  $\alpha_C^{(1)} = 2/3$ ,  $\alpha_U^{(1)} = 0$  respectively; with probability  $1/2$  we are in setting  $Y$ , the fractions of agents with types  $F, C, U$  are  $\alpha_F^{(2)} = 0$ ,  $\alpha_C^{(2)} = 1$ ,  $\alpha_U^{(2)} = 0$  respectively. This finishes the description of the instance. Note there are 2 worlds and 2 settings yielding 4 possible environments which we label  $LX$ ,  $LY$ ,  $HX$ , and  $HY$ .

Let  $\tau = 0.001$ . Suppose there exists a mechanism  $\mathcal{M}$  that outputs the majority wish with probability at least  $1 - \tau$  in a symmetric strategy profile  $\Sigma$ . We will show that  $\Sigma$  cannot be a  $\tau$ -strong Bayes Nash equilibrium.

Notice that the contingent agents are the majority in all the four settings. In both environments  $HX$  and  $HY$ , each of which happens with probability 1%, the majority wish is always to accept. Thus, the mechanism must accept with probability at least 99% in each environment for otherwise it will be far from being achieving  $1 - \tau$  accuracy. Similarly, in both environments  $LX$  and  $LY$ , each of which occurs with probability 49%, the mechanism must accept with probability at most 1% for otherwise it will be far from being achieving  $1 - \tau$  accuracy.

To show that  $\Sigma$  cannot be a  $\tau$ -strong Bayes Nash equilibrium, consider that the set of deviating agents are all the candidate-friendly agents. Those candidate-friendly agents pretend they are contingent agents such that signal  $\ell$  is received with probability  $1/2$  and signal  $h$  with probability  $1/2$ . These candidate-friendly agents will follow the strategy of the real contingent agents according to  $\Sigma$ . Let  $\Sigma'$  be the resultant strategy profile. We aim to show that those deviating candidate-friendly agents can increase their utilities significantly in  $\Sigma'$ .

For an intuitive argument, in  $LX$ , the mechanism sees that all the agents are contingent, and the fraction of agents receiving signal  $h$  follows distribution  $\text{Bin}(T/3, 1/2)$ ; in  $HY$ , the mechanism also sees that all the agents are contingent, and the fraction of agents receiving signal  $h$  follows distribution  $\text{Bin}(T, 1/6)$ . Lemma D.5 implies that the mechanism cannot distinguish between environments  $LX$  and  $HY$  with probability more than 87.7%. Before deviating, the mechanism will output  $\mathbf{A}$  with probability at most 1% in  $LX$ ; after deviating, the mechanism will output  $\mathbf{A}$  with probability at least  $99\% \cdot 87.7\% > 1\%$  in  $LX$  by confusing  $LX$  with  $HY$ . The candidate-friendly agents will benefit from deviating.

To make the arguments in the previous paragraph more rigorous, let  $\mathcal{M}(X) = 1$  if the mechanism outputs  $\mathbf{A}$  when all the agents are contingent, agents play according to  $\Sigma$ , and the number of agents receiving signal  $h$  is  $X$ . Let  $\mathcal{M}(X) = 0$  if the output is  $\mathbf{R}$  under the same circumstance. This definition is well-defined due to Proposition D.4.

In  $HY$ , the mechanism outputs  $\mathbf{A}$  with probability  $\mathbb{E}_{X \sim \text{Bin}(T, 1/6)}[\mathcal{M}(X)]$ . In  $LX$ , when the candidate-friendly agents deviate to  $\Sigma'$ , the mechanism outputs  $\mathbf{A}$  with probability  $\mathbb{E}_{X \sim \text{Bin}(T/3, 1/2)}[\mathcal{M}(X)]$ . Lemma D.5 implies

$$\left| \mathbb{E}_{X \sim \text{Bin}(T, 1/6)}[\mathcal{M}(X)] - \mathbb{E}_{X \sim \text{Bin}(T/3, 1/2)}[\mathcal{M}(X)] \right| < 0.123.$$

Since we have shown that the mechanism outputs  $\mathbf{A}$  with probability at least 99% in  $HY$ , the mechanism outputs  $\mathbf{A}$  with probability at least 86.7% in  $LX$  when the candidate-friendly agents deviate. Since environment  $LX$  happens with probability 0.49, the expected utility for each candidate-friendly agent  $t$  is at least  $0.49 \times 86.7\% \times v_t(L, \mathbf{A}) = 0.84966$ . However, without deviating,  $\mathbf{A}$  will

be output with probability at most  $0.98 \cdot 1\% + 0.02 = 0.0298$ , and the expected utility for each candidate-friendly agent  $t$  is upper-bounded by  $0.0298 \times v_t(H, \mathbf{A}) = 0.0894$ . We have seen that the candidate-friendly agents receive a utility gain of at least  $0.76026 > \tau$ .  $\square$

As a remark, our impossibility result Theorem D.2 holds even for randomized mechanism. If mechanism can be randomized, Proposition D.4 becomes that  $\{\alpha_F, \alpha_U, \alpha_C\} \cup \{\beta_{F,\ell}, \beta_{F,h}, \beta_{U,\ell}, \beta_{U,h}, \beta_{C,\ell}, \beta_{C,h}\}$  completely determines the *probability* that the mechanism output  $\mathbf{A}$  (or  $\mathbf{R}$ ). In the proof of Theorem D.2,  $\mathcal{M}(X)$  becomes the *probability* that the mechanism output  $\mathbf{A}$ , rather than either 0 or 1. The remaining part of the proof is exactly the same.

## E Extension to Non-binary Worlds and Signals

In Sect. E.1, we generalized the model to the setting with non-binary worlds and non-binary signals. In Sect. E.2, we present our mechanism for the setting with binary signals and non-binary worlds. In Sect. E.3, we show that the generalization to non-binary signals is straightforward.

### E.1 Model and Preliminaries

In our non-binary model, as in our binary model,  $T$  agents are voting for two *alternatives*,  $\mathbf{A}$  and  $\mathbf{R}$  (corresponding to “accept” and “reject”). However, in our non-binary model there is a set of  $N$  possible *worlds* (or *states*)  $\mathcal{W} = \{1, \dots, N\}$ , where the higher the value the more  $\mathbf{A}$  is preferred to  $\mathbf{R}$ . Agents do not know which world is the actual world that they are in. They have a prior common belief on the likelihood of each world. Let  $W$  be the actual world which is viewed as a random variable. Let  $P_n = \Pr(W = n)$  be the prior over worlds. Each agent knows the values of  $P_1, \dots, P_n$  as prior beliefs. We further assume  $P_n > 0$  for each  $n$ , for otherwise we can remove world  $n$  from  $\mathcal{W}$  without loss of generality.

Each agent will then receives a *signal* from the set  $\mathcal{S} = \{1, \dots, M\}$ . Let  $S_t$  be the random variable representing the signal that agent  $t$  receives. Given  $W = n$ , for any  $n$ , the signals agents receive have the same distribution and are conditionally independent. Let  $P_{mn} = \Pr(S_t = m \mid W = n)$  be the probability that signal  $m$  will be received (by an arbitrary agent  $t$ ) if the actual world is  $n$ . The set of values  $\{P_{mn} : m = 1, \dots, M; n = 1, \dots, N\}$  is known by all the agents. Signals are positively correlated to the worlds:

$$\Pr(S_t \geq m \mid W = n_1) = \sum_{m'=m}^M P_{m'n_1} > \Pr(S_t \geq m \mid W = n_2) = \sum_{m'=m}^M P_{m'n_2} \quad (13)$$

for any worlds  $n_1 > n_2$ , any signal  $m$ , and any agent  $t$ .

The remaining definitions for the non-binary model in this section are rather analogous to the binary case, but we include them for completeness.

Each agent  $t$  has a *utility function*  $v_t : \mathcal{W} \times \{\mathbf{A}, \mathbf{R}\} \rightarrow \{0, 1, \dots, B\}$ . As mentioned earlier, a higher value of  $W$  indicates  $\mathbf{A}$  is more preferable:  $v_t(n_1, \mathbf{A}) > v_t(n_2, \mathbf{A})$  and  $v_t(n_1, \mathbf{R}) < v_t(n_2, \mathbf{R})$  for any  $n_1, n_2 \in \mathcal{W}$  with  $n_1 > n_2$ . Since we can always rescale agents’ utilities, for simplicity, we assume without loss of generality that agents’ utilities are integers and bounded by  $B \in \mathbb{Z}^+$ . Agents, with their prior beliefs and receiving signals, will have posterior beliefs about the distribution of  $W$  and react to the mechanism in a way maximizing their expected utilities.

We assume  $v_t(n, \mathbf{A}) \neq v_t(n, \mathbf{R})$  for each agent  $t$  and each  $n \in \mathcal{W}$ , so that agents always strictly prefer one alternative over the other. Given a world  $n$ , let  $T(\mathbf{A}, n) = \{t \mid v_t(\mathbf{A}, n) > v_t(\mathbf{R}, n)\}$  be the set of agents that prefer  $\mathbf{A}$  in world  $n$  and let  $\alpha_n^{\mathbf{A}} = \frac{|T(\mathbf{A}, n)|}{|T|}$  be the fraction of agents that prefer  $\mathbf{A}$  in world  $n$ . We can similarly define  $T(\mathbf{R}, n)$  and  $\alpha_n^{\mathbf{R}} = \frac{|T(\mathbf{R}, n)|}{|T|} = 1 - \alpha_n^{\mathbf{A}}$ . Naturally,  $\alpha_n^{\mathbf{A}}$  is increasing in  $n$  (when the underlying quality of the candidate increases, more agents prefer  $\mathbf{A}$ ) and  $\alpha_n^{\mathbf{R}}$  is decreasing in  $n$ . As before, we assume that the  $\alpha_n^{\mathbf{A}}$  and  $\alpha_n^{\mathbf{R}}$  are common knowledge, which is natural in many scenarios, including the faculty candidate hiring example. If this assumption does not hold, results in Sect. D show that we cannot achieve the truthful guarantee even under the setting  $M = N = 2$ .



For any world  $n$ , let

$$M(n) = \begin{cases} \mathbf{A} & \alpha_n^{\mathbf{A}} > \frac{1}{2} \\ \mathbf{R} & \text{otherwise} \end{cases}$$

be the majority preference if the actual world were  $n$ . We assume that  $T$  is an odd number to avoid ties.

**Definition E.1.** Given a utility profile  $\{v_1, \dots, v_T\}$  and letting  $n^*$  be the actual world, we say  $M(n^*)$  is the *majority wish*.

The goal is to output the majority wish  $M(n^*)$ , the alternative that is preferred by at least half of the agents in the actual world.

Our results will sometimes require  $T$ , the number of agents, to be sufficiently large, and it may be helpful to think of  $T \rightarrow \infty$ . However, we will always assume that the parameters of the model:  $B$ ,  $\{P_n\}_{n \in \mathcal{W}}$ ,  $\{P_{mn}\}_{m \in \mathcal{S}, n \in \mathcal{W}}$ , and  $\{\alpha_n^{\mathbf{A}}, \alpha_n^{\mathbf{R}}\}_{n \in \mathcal{W}}$ , do not depend on  $T$  in any way.

In the faculty candidate hiring example, the worlds  $\mathcal{W} = \{1, \dots, N\}$  describe the quality of the candidate, with 1 being the worst and  $N$  being the best. The signals  $S_t \in \{1, \dots, M\}$  correspond to the impression of this candidate, with  $S_t = 1$  being the worst impression and  $S_t = M$  being the best impression. It is natural to assume that  $S_t$ 's are positively correlated to  $W$ , which agrees with our model.

### E.1.1 Candidate-Friendly, Contingent and Candidate-Unfriendly Agents

Let  $\mathcal{L} = \{n \in \mathcal{W} \mid \alpha_n^{\mathbf{A}} < \frac{1}{2}\}$  and  $\mathcal{H} = \{n \in \mathcal{W} \mid \alpha_n^{\mathbf{A}} > \frac{1}{2}\}$ . Since we cannot have  $\alpha_n^{\mathbf{A}} = \frac{1}{2}$  for an odd  $T$ ,  $\{\mathcal{L}, \mathcal{H}\}$  is a partition of  $\mathcal{W}$ . In addition, since  $\alpha_n^{\mathbf{A}}$  is increasing in  $n$ , there exists a threshold such that all those  $n$  below the threshold belong to  $\mathcal{L}$  and all those  $n$  above belong to  $\mathcal{H}$ . Indeed,  $\mathcal{L}$  is the set of ‘‘low quality’’ worlds where  $\mathbf{R}$  is preferred, and  $\mathcal{H}$  is the set of ‘‘high quality’’ world is preferred.

For each agent  $t$ , define  $\mathcal{L}_t = \{n \in \mathcal{W} \mid v_t(n, \mathbf{R}) > v_t(n, \mathbf{A})\}$  and  $\mathcal{H}_t = \{n \in \mathcal{W} \mid v_t(n, \mathbf{A}) > v_t(n, \mathbf{R})\}$ . Then  $\mathcal{L}_t$  is the set of those low quality worlds based on agent  $t$ 's utility where  $\mathbf{R}$  is preferred, and  $\mathcal{H}_t$  is the set of those high quality worlds where  $\mathbf{A}$  is preferred for  $t$ . Since  $v_t(n, \mathbf{A}) - v_t(n, \mathbf{R})$  is increasing in  $n$  (the first term is increasing and the second term is decreasing), each agent  $t$  also has a personal threshold that separate  $\mathcal{W}$  to  $\mathcal{L}_t, \mathcal{H}_t$ . We can define the candidate-friendly agents, contingent agents and candidate-unfriendly agents based on whether the personal threshold is below, equal to, or above the average threshold.

We say an agent  $t$  is candidate-friendly if  $\mathcal{H}_t \cap \mathcal{L} \neq \emptyset$ . This says that there exists a world  $n \in \mathcal{W}$  where the fraction of agents preferring  $\mathbf{A}$  is below  $1/2$  (i.e.,  $\alpha_n^{\mathbf{A}} < \frac{1}{2}$ ) but  $t$  still prefers  $\mathbf{A}$ . Equivalently, an agent  $t$  is candidate-friendly if  $\mathcal{L}_t \subsetneq \mathcal{L}$ , or  $\mathcal{H} \subsetneq \mathcal{H}_t$ . Correspondingly, an agent  $t$  is candidate-unfriendly if  $\mathcal{L}_t \cap \mathcal{H} \neq \emptyset$ , or equivalently,  $\mathcal{L} \subsetneq \mathcal{L}_t$ , or  $\mathcal{H}_t \subsetneq \mathcal{H}$ . An agent  $t$  is contingent if  $\mathcal{L}_t = \mathcal{L}$ , or equivalently,  $\mathcal{H}_t = \mathcal{H}$ . We still use  $F, C, U$  to denote the three types of agents, and use  $\alpha_F, \alpha_C, \alpha_U$  to denote their fractions as before. As a remark, a candidate-friendly agent (candidate-unfriendly agent resp.) does not always prefer  $\mathbf{A}$  ( $\mathbf{R}$  resp.) like before, (s)he merely has a threshold below (above resp.) the average.

Let  $L = \max\{n \in \mathcal{L}\}$  be the maximum world where  $\mathbf{R}$  is preferred by the majority, and  $H = \min\{n \in \mathcal{H}\}$  be the minimum world where  $\mathbf{A}$  is preferred by the majority. We clearly have  $H = L + 1$ . For each agent  $t$ , let  $L_t = \max\{n \in \mathcal{L}_t\}$  be the maximum world where  $\mathbf{R}$  is preferred, and let  $H_t = \min\{n \in \mathcal{H}_t\}$  be the minimum world where  $\mathbf{A}$  is preferred. Set  $L_t = 0$  if  $\mathcal{L}_t = \emptyset$  and  $H_t = N + 1$  if  $\mathcal{H}_t = \emptyset$ . Clearly,  $H_t = L_t + 1$ .

### E.1.2 Additional Notations

Given a strategy profile  $\Sigma = \{\sigma_1, \dots, \sigma_T\}$  and a mechanism  $\mathcal{M}$ , let  $\lambda_n^{\mathbf{A}, \mathcal{M}}(\Sigma)$  be the probability that alternative  $\mathbf{A}$  is announced as the winner given the actual world is  $n$ , then  $\lambda_n^{\mathbf{R}, \mathcal{M}}(\Sigma) = 1 - \lambda_n^{\mathbf{A}, \mathcal{M}}(\Sigma)$  is the probability that alternative  $\mathbf{R}$  wins given the actual world is  $n$ . We will omit the superscript  $\mathcal{M}$  when it is clear what mechanism we are discussing.

All the agents' *ex-ante* utilities depend exclusively on  $\lambda_1^{\mathbf{A}}(\Sigma), \dots, \lambda_N^{\mathbf{A}}(\Sigma)$  (or  $\lambda_1^{\mathbf{R}}(\Sigma), \dots, \lambda_N^{\mathbf{R}}(\Sigma)$ ), and each agent  $t$ 's utility is given by

$$u_t(\Sigma) = \sum_{n=1}^N P_n (\lambda_n^{\mathbf{A}}(\Sigma)v_t(n, \mathbf{A}) + \lambda_n^{\mathbf{R}}(\Sigma)v_t(n, \mathbf{R})) \quad (14)$$

By substituting  $\lambda_n^{\mathbf{R}}(\Sigma) = 1 - \lambda_n^{\mathbf{A}}(\Sigma)$ ,

$$u_t(\Sigma) = \sum_{n=1}^N P_n v_t(n, \mathbf{R}) + \sum_{n=1}^N P_n \lambda_n^{\mathbf{A}}(\Sigma)(v_t(n, \mathbf{A}) - v_t(n, \mathbf{R})). \quad (15)$$

We will always use  $\Sigma^* = \{\sigma_1^*, \dots, \sigma_T^*\}$  to denote the truthful strategy profile.

Table 1 lists all the frequently used notations.

notation	meaning
$\mathcal{W} = \{1, \dots, N\}$	the set of all worlds
$\mathcal{S} = \{1, \dots, M\}$	the set of all signals
$P_n$	the prior belief for the probability the actual world is $n$
$P_{mn}$	the probability of receiving signal $m$ under world $n$
$v_t(n, \mathbf{A}), v_t(n, \mathbf{R})$	the ( <i>ex-post</i> ) utility for agent $t$ for alternative $\mathbf{A}, \mathbf{R}$ if the actual world is $n$
$u_t(\Sigma)$	the ( <i>ex-ante</i> ) expected utility for agent $t$ given strategy profile $\Sigma$
$\alpha_n^{\mathbf{A}}, \alpha_n^{\mathbf{R}}$	the fraction of agents preferring $\mathbf{A}, \mathbf{R}$ under world $n$
$M(n)$	the majority favored alternative under world $n$
$\mathcal{L}$	the set of worlds where more than half of the agents prefer $\mathbf{R}$
$\mathcal{H}$	the set of worlds where more than half of the agents prefer $\mathbf{A}$
$\mathcal{L}_t$	the set of worlds where $\mathbf{R}$ is preferred for agent $t$
$\mathcal{H}_t$	the set of worlds where $\mathbf{A}$ is preferred for agent $t$
$L$	the maximum world where $\mathbf{R}$ is preferred by the majority
$H$	the minimum world where $\mathbf{A}$ is preferred by the majority
$L_t$	the maximum world where $\mathbf{R}$ is preferred for agent $t$
$H_t$	the minimum world where $\mathbf{A}$ is preferred for agent $t$
$F, C, U$	candidate-friendly agents, contingent agents, candidate-unfriendly agents
$\alpha_F, \alpha_C, \alpha_U$	fractions of the three types of agents
$\lambda_n^{\mathbf{A}}(\Sigma), \lambda_n^{\mathbf{R}}(\Sigma)$	the probability a given mechanism outputs $\mathbf{A}, \mathbf{R}$ for strategy profile $\Sigma$
$\Sigma^*$	the truthful strategy profile

Table 1: Table of notations.

## E.2 Non-binary Worlds

In this section, we consider the generalization to the setting with more than two worlds  $N > 2$ , while keeping the binary signal assumption  $M = 2$ . We will see in the next section that the generalization to non-binary signals is simple. For this section, we will again use  $\ell$  to denote signal 1 and  $h$  to denote signal 2.

In the case  $M = N = 2$ , we have asked each agent his/her received signal, type, and posterior belief on the fraction of agents who will report signal 1. In the case  $N > 2$  here, while it is still natural to ask an agent for his/her signal, asking for a posterior prediction and keeping the mechanism as before will not work here. In particular, this will make Theorem 3.2 fail. To reason this intuitively, it is easy to see that, if the mechanism is required to output the alternative favored by the majority, the mechanism will output  $\mathbf{A}$  if the actual world is in  $\mathcal{H}$  and output  $\mathbf{R}$  if the actual world is in  $\mathcal{L}$ . To ensure this, we need to make sure the median of the posterior prediction is between  $P_{hL}$  and  $P_{hH}$ . However, while this is true for  $N = 2$  as Theorem 3.1 suggests (in fact, all the possible posterior predictions,  $T_{h\ell}$  and  $T_{hh}$ , are between  $P_{hL}$  and  $P_{hH}$ ), this is not necessarily true for  $N > 2$ .

As a solution to this issue, for each contingent agent  $t$ , we ask him/her for a value between  $P_{hL_t}$  and  $P_{hH_t}$  (agents with  $\mathcal{L}_t = \emptyset$  report a value between 0 and  $P_{h1}$  and agents with  $\mathcal{H}_t = \emptyset$  report a value between  $P_{hN}$  and 1), and the median of these values will be between  $P_{hL}$  and  $P_{hH}$ . A natural way to ask an agent for this value can be, *please give a percentage value  $q$  such that you would like*

alternative **A** if the fraction of agents reporting signal  $h$  is more than  $q$  percent, and you would like alternative **R** otherwise.

Our mechanism is presented in Mechanism 3.<sup>7</sup>

---

**Mechanism 3** The Wisdom-of-the-Crowd-Voting Mechanism for  $N > 2$  and  $M = 2$

---

- 1: Each agent  $t$  reports to the mechanism the signal (s)he receives (either  $\ell$  or  $h$ ), the type (either  $F$ ,  $C$ , or  $U$ ).
  - 2: If agent  $t$  reports type  $F$ , his/her reported signal will be automatically treated as  $h$ ; if agent  $t$  reports type  $U$ , his reported signal will be automatically treated as  $\ell$ .
  - 3: If an agent reports type  $C$ , ask him/her to report a value  $q_t \in [0, 1]$  such that (s)he would like **A** if and only if the fraction of agents who report  $h$  is more than  $q_t$ . The value  $q_t$  should be given with the treatment in the previous step being considered, and the mechanism makes this clear to the agents. For an agent with type  $F$ , set  $q_t = 0$ . For an agent with type  $U$ , set  $q_t = 1$ .
  - 4: Compute the median of those  $q_t$ , denoted by  $\bar{q}$ .
  - 5: If the fraction of the agents reporting  $h$  is more than the median  $\bar{q}$ , announce **A** being the winning alternative; otherwise, announce **R** being the winning alternative.
- 

In our faculty candidate hiring example, the questionnaire corresponding to the mechanism looks like the following:

1. Choose one of the following: as compared with the average faculty member, independent of the candidate's qualification:
  - I I am more predisposed toward rejection;
  - II I am more predisposed toward acceptance;
  - III Neither I nor II, i.e., I am with the average faculty member.
2. What is your impression of this candidate during the individual interview between you and this candidate?
  - (a) I had a good impression.
  - (b) I did not have a good impression.
3. Your "provisional ballot" will be cast as follows:
  - Accept: if you choose II in Question 1, or, if you choose III in Question 1 and (a) in Question 2;
  - Reject: if you choose I in Question 1, or, if you choose III in Question 1 and (b) in Question 2.
4. If you answered III for Question 1, what fraction of provisional ballots do you predict will be cast for accept?

Next, we will prove that the mechanism outputs the alternative favored by the majority with high probability, and the truthful strategy profile is an  $\varepsilon$ -strong Bayes Nash equilibrium for  $\varepsilon = o(1)$ . For simplicity and clarity in describing the ideas behind the proofs, we will not perform the Chernoff bound analyses as in Section B.3, and we will assume  $T \rightarrow \infty$  as it is in Sect. 3. As a result, for any world  $n$ , the fractions of agents receiving signal  $\ell$  and  $h$  are, almost surely,  $P_{\ell n}$  and  $P_{hn}$  respectively.

**Theorem E.2.** *Suppose  $T \rightarrow \infty$ . If all the agents play the truthful strategy, then our mechanism outputs an alternative favored by more than half of the agents.*

*Proof.* Suppose all the agents report truthfully. Let  $n$  be the actual world. We need to show that **A** is announced if and only if  $n \in \mathcal{H}$ . We assume  $n \in \mathcal{H}$  without loss of generality, as the analysis for  $n \in \mathcal{L}$  is similar. When all the agents play the truthful strategy profile  $\Sigma^*$ , the fraction of agents reporting signal  $h$  is  $\alpha_F + \alpha_C \cdot P_{hn}$ .

On the other hand, for any agent  $t$ , (s)he prefers **A** if  $n \in \mathcal{H}_t$ , and (s)he prefers **R** if  $n \in \mathcal{L}_t$ . If (s)he were asked to give a value such that (s)he would like **A** if and only if the fraction of agents who receive signal  $h$  is less than this value, (s)he would have report a value between  $P_{hL_t}$

---

<sup>7</sup>We can consider the same modification as in the binary case (Section C).

and  $P_{hH_t}$ . Now, considering that, as instructed by the mechanism, those  $\alpha_F \cdot T$  (resp.  $\alpha_U \cdot T$ ) agents will always *report* signal  $h$  (resp. signal  $\ell$ ) regardless of what they receive, (s)he will report  $q_t \in (\alpha_F + \alpha_C \cdot P_{hL_t}, \alpha_F + \alpha_C \cdot P_{hH_t})$ .

It is then easy to see that the median  $\bar{q}$  is between  $\alpha_F + \alpha_C \cdot P_{hL}$  and  $\alpha_F + \alpha_C \cdot P_{hH}$ , which is less than the fraction of agents reporting signal  $h$  (which is  $\alpha_F + \alpha_C \cdot P_{hn}$  as computed earlier). The last step of our mechanism will make sure **A** is output.  $\square$

As a remark, if we do not assume  $T \rightarrow \infty$ , to show that the statement in Theorem E.2 fails with an exponentially low probability, we need to make an extra assumption that the median  $\bar{q}$  is not exponentially close to the two endpoints  $\alpha_F + \alpha_C \cdot P_{hL}$  and  $\alpha_F + \alpha_C \cdot P_{hH}$ . This is a natural assumption, as an agent's reported value should not depend on  $T$ . In addition, in practice, it is natural to expect that most agents will report values that are around the midpoint of the interval  $(\alpha_F + \alpha_C \cdot P_{hL_t}, \alpha_F + \alpha_C \cdot P_{hH_t})$ .

Next, we show that our mechanism satisfies the truthful property. Again, we consider  $T \rightarrow \infty$ .

**Theorem E.3.** *Suppose  $T \rightarrow \infty$ . The truthful strategy profile form a strong Bayes Nash Equilibrium.*

The ideas behind the proof of Theorem E.3 is similar as before. Firstly, those contingent agents will not deviate from the truthful strategy, as their utilities have already been maximized. Secondly, we can prove a lemma similar to Lemma 3.5 showing that there is a conflict of interest between an arbitrary candidate-friendly agent and an arbitrary candidate-unfriendly agent. This shows that the set of deviating agents  $D$  can only contain either candidate-friendly agents or candidate-unfriendly agents. This further implies that more than half of the agents are truth-telling. Finally, the use of median in our mechanism ensures that less than half of the agents' deviating cannot change the output alternative.

**Lemma E.4.** *Let  $\Sigma^*$  be the truthful strategy profile and  $\Sigma'$  be an arbitrary strategy profile. Let  $t_1$  be an arbitrary candidate-friendly agent and  $t_2$  be an arbitrary candidate-unfriendly agent. Suppose  $T \rightarrow \infty$ . We have*

1. *If  $u_{t_1}(\Sigma') > u_{t_1}(\Sigma^*)$ , then  $u_{t_2}(\Sigma') < u_{t_2}(\Sigma^*)$ .*
2. *If  $u_{t_1}(\Sigma') < u_{t_1}(\Sigma^*)$ , then  $u_{t_2}(\Sigma') > u_{t_2}(\Sigma^*)$*

The proof of this lemma is more involved than that of Lemma 3.5, but the ideas behind are similar.

*Proof.* We will only prove (1), as the proof for (2) is similar. For ease of notation, in this proof, we assume  $t_1 = 1$  and  $t_2 = 2$  without loss of generality. Suppose  $u_1(\Sigma') > u_1(\Sigma^*)$ , and we aim to show  $u_2(\Sigma') < u_2(\Sigma^*)$ . Theorem E.2 implies that  $\lambda_n(\Sigma^*) = 0$  for all  $n \in \mathcal{L}$  and  $\lambda_n(\Sigma^*) = 1$  for all  $n \in \mathcal{H}$ . Firstly, we show that

$$\sum_{n \in \mathcal{L} \cap \mathcal{H}_1} P_n (\lambda_n^A(\Sigma') - \lambda_n^A(\Sigma^*)) > \sum_{n \in \mathcal{H} \cap \mathcal{L}_2} P_n (\lambda_n^A(\Sigma^*) - \lambda_n^A(\Sigma')). \quad (16)$$

This is because

$$\begin{aligned}
0 &< u_1(\Sigma') - u_1(\Sigma^*) && \text{(by our assumption)} \\
&= \sum_{n=1}^N P_n (\lambda_n^{\mathbf{A}}(\Sigma') - \lambda_n^{\mathbf{A}}(\Sigma^*)) (v_1(n, \mathbf{A}) - v_1(n, \mathbf{R})) && \text{(by (15))} \\
&\leq \sum_{n \in \mathcal{W} \setminus (\mathcal{L}_1 \cup \mathcal{H}_2)} P_n (\lambda_n^{\mathbf{A}}(\Sigma') - \lambda_n^{\mathbf{A}}(\Sigma^*)) (v_1(n, \mathbf{A}) - v_1(n, \mathbf{R})) && (\dagger) \\
&= \sum_{n \in \mathcal{L} \cap \mathcal{H}_1} P_n (\lambda_n^{\mathbf{A}}(\Sigma') - \lambda_n^{\mathbf{A}}(\Sigma^*)) (v_1(n, \mathbf{A}) - v_1(n, \mathbf{R})) \\
&\quad - \sum_{n \in \mathcal{H} \cap \mathcal{L}_2} P_n (\lambda_n^{\mathbf{A}}(\Sigma^*) - \lambda_n^{\mathbf{A}}(\Sigma')) (v_1(n, \mathbf{A}) - v_1(n, \mathbf{R})) \\
&\leq \sum_{n \in \mathcal{L} \cap \mathcal{H}_1} P_n (\lambda_n^{\mathbf{A}}(\Sigma') - \lambda_n^{\mathbf{A}}(\Sigma^*)) (v_1(L, \mathbf{A}) - v_1(L, \mathbf{R})) \\
&\quad - \sum_{n \in \mathcal{H} \cap \mathcal{L}_2} P_n (\lambda_n^{\mathbf{A}}(\Sigma^*) - \lambda_n^{\mathbf{A}}(\Sigma')) (v_1(L, \mathbf{A}) - v_1(L, \mathbf{R})), && (\ddagger)
\end{aligned}$$

which implies (16), where both Step  $(\dagger)$  and  $(\ddagger)$  are based on the following facts. In particular,  $(\dagger)$  is based on the first two facts, and  $(\ddagger)$  is based on the first and the third facts.

- for  $n \in \mathcal{L}$ ,  $\lambda_n^{\mathbf{A}}(\Sigma') - \lambda_n^{\mathbf{A}}(\Sigma^*) = \lambda_n^{\mathbf{A}}(\Sigma') - 0 \geq 0$ ; for  $n \in \mathcal{H}$ ,  $\lambda_n^{\mathbf{A}}(\Sigma^*) - \lambda_n^{\mathbf{A}}(\Sigma') = 1 - \lambda_n^{\mathbf{A}}(\Sigma') \geq 0$ .
- $v_1(n, \mathbf{A}) - v_1(n, \mathbf{R})$  is negative for  $n \in \mathcal{L}_1$  and is positive for  $n \in \mathcal{H}_2$ . Notice that this is also true for  $v_2$ .
- the expression  $v_1(n, \mathbf{A}) - v_1(n, \mathbf{R})$  is increasing in  $n$ . This is true for any agent  $t$ . In particular, for each agent  $t$ ,  $v_t(n, \mathbf{A})$  is increasing in  $n$  and  $v_t(n, \mathbf{R})$  is decreasing in  $n$ .

Next, we show that (16) implies  $u_2(\Sigma') < u_2(\Sigma^*)$ . By the same calculations and analyses above, we have

$$\begin{aligned}
u_2(\Sigma') - u_2(\Sigma^*) &\leq \sum_{n \in \mathcal{L} \cap \mathcal{H}_1} P_n (\lambda_n^{\mathbf{A}}(\Sigma') - \lambda_n^{\mathbf{A}}(\Sigma^*)) (v_2(L, \mathbf{A}) - v_2(L, \mathbf{R})) \\
&\quad - \sum_{n \in \mathcal{H} \cap \mathcal{L}_2} P_n (\lambda_n^{\mathbf{A}}(\Sigma^*) - \lambda_n^{\mathbf{A}}(\Sigma')) (v_2(L, \mathbf{A}) - v_2(L, \mathbf{R})) \\
&\hspace{15em} \text{(same calculations above)} \\
&< 0, && \text{(by } v_2(L, \mathbf{A}) - v_2(L, \mathbf{R}) < 0 \text{ and (16))}
\end{aligned}$$

which implies the lemma.  $\square$

Now we are ready to prove Theorem E.3.

*Proof of Theorem E.3.* Suppose otherwise and there is a set of deviating agents  $D$ . Let  $\Sigma'$  be the profile after the deviation of agents in  $D$ . Firstly, we show that  $D$  cannot contain a contingent agent. Notice that such an agent's utility has already been maximized by the truthful profile  $\Sigma^*$ . Suppose  $\lambda_n^{\mathbf{A}}(\Sigma') \neq \lambda_n^{\mathbf{A}}(\Sigma^*)$  for certain  $n$ . It must be that  $\lambda_n^{\mathbf{A}}(\Sigma') > \lambda_n^{\mathbf{A}}(\Sigma^*) = 0$  if  $n \in \mathcal{L}$ , and  $\lambda_n^{\mathbf{A}}(\Sigma') < \lambda_n^{\mathbf{A}}(\Sigma^*) = 1$  if  $n \in \mathcal{H}$ . It is then easy to see that this agent's utility will decrease, which contradicts to 2 of Definition 2.1. Suppose  $\lambda_n^{\mathbf{A}}(\Sigma') = \lambda_n^{\mathbf{A}}(\Sigma^*)$  for all  $n$ . We have  $u_t(\Sigma') = u_t(\Sigma^*)$  for every agent  $t$ . This already contradicts to 3 of Definition 2.1.

Next, Lemma E.4 ensures that  $D$  cannot contain both a candidate-friendly agent and a candidate-unfriendly agent. Assume without loss of generality that  $D$  only contains candidate-friendly agents. In order to maximize the chance that  $\mathbf{A}$  is output, those candidate-friendly agents need to maximize the fraction of agents reporting signal  $h$  and minimize the median  $\bar{q}$ . However, the mechanism always does this for them in the truth-telling profile  $\Sigma^*$ : Step 2 makes sure they report signal  $h$ , and Step 3 makes sure they report  $q_t = 0$ . Therefore, those candidate-friendly agents' utilities are maximized by truth-telling, which contradicts to our assumption for  $D$ .  $\square$

### E.3 Non-binary Signals

There is a simple reduction from the non-binary signal setting to the binary-signal setting. Suppose the signal space is  $\{1, \dots, M\}$ . To reduce it to a binary signal space  $\{\ell, h\}$ , we set an arbitrary non-integer number  $s_\top$  between 1 and  $M$ . All the signals less than  $s_\top$  are reduced to the “bad” signal  $\ell$ , and all the signals greater than  $s_\top$  are reduced to the “good” signal  $h$ . The mechanisms in the previous sections can be adapted to the setting here. The mechanisms are the same as before, except for the following change: whenever the mechanism asks an agent for a binary signal in the previous setting, the mechanism asks the agent whether the signal (s)he received is less than or more than  $s_\top$ , which corresponds to signal  $\ell$  and  $h$  respectively.

For the mechanism in Section E.2, all the properties, including that the mechanism outputs the alternative favored by more than half of the agents and that the truth-telling strategy profile form a strong Bayes Nash Equilibrium, continue to hold in the non-binary signal setting with exactly the same proofs.

If we are dealing with binary world  $\mathcal{W} = \{L, H\}$ , for the mechanism in Section 3, it is easy to see that these properties also continue to hold here if we prove the following inequality that is similar to the one in Theorem 3.1:

$$\forall m \in \{1, \dots, M\} : P_{hH} > T_{hm} > P_{hL}, \quad (17)$$

where  $P_{hL}$  and  $P_{hH}$  are the probabilities that a signal above  $s_\top$  is received if the actual world is  $L$  and  $H$  respectively, and  $T_{hm}$  is the probability that an agent who receives signal  $m$  believes that another agent will receive a signal that is more than  $s_\top$ . Intuitively, if (17) holds, all the agents’ posterior predictions are still between  $P_{hL}$  and  $P_{hH}$ , and the majority wish will still be “surprisingly popular”. The proof of (17) is by straightforward Bayesian analysis, and is left to the readers.