

Making Consensus Tractable

Elchanan Mossel* and Omer Tamuz†

October 7, 2013

Abstract

We study a model of consensus decision making, in which a finite group of Bayesian agents has to choose between one of two courses of action. Each member of the group has a private and independent signal at his or her disposal, giving some indication as to which action is optimal. To come to a common decision, the participants perform repeated rounds of voting. In each round, each agent casts a vote in favor of one of the two courses of action, reflecting his or her current belief, and observes the votes of the rest.

We provide an efficient algorithm for the calculation the agents have to perform, and show that consensus is always reached and that the probability of reaching a wrong decision decays exponentially with the number of agents.

1 Introduction

Consensus voting, or decision by unanimous agreement, is a method of communal governance that requires all members of a group to agree on a chosen course of action. The European Union’s Treaty of Lisbon [6] decrees that “*Except where the Treaties provide otherwise, decisions of the European Council shall be taken by consensus.*” In this the EU follows the historical example of the Diet of the Hanseatic League [14] and others.

Proponents of this method consider it to have many advantages over majority voting: it cultivates discussion, participation and responsibility, and avoids the so-called “tyranny of the majority”. The drawback is, of course, a lengthy and difficult decision making process, lacking even the guarantee of a conclusive ending.

However, in standard theoretical setups of rational Bayesian participants (e.g. [13], [7]), agents cannot “agree to disagree” [2], and consensus is eventually reached. Unfortunately, this may come at the price of tractability; Bayesian calculations can, in some situations, be practically impossible [7].

Indeed, modeling economic behavior involves an inherent conflict between rationality and tractability [20]. It seems that in many situations it is practically impossible to calculate which course of action is optimal, leaving the theoretician with a model that is either not rational, and thus hard to justify, or not tractable and hence not realistic. A common course of action is to relax the rationality assumption and consider boundedly-rational agents. We do not do this, but rather show that a fully rational model is, perhaps unexpectedly, tractable.

*UC Berkeley and Weizmann Institute of Science. Supported by a Sloan fellowship in Mathematics, by BSF grant 2004105, by NSF Career Award (DMS 054829) by ONR award N00014-07-1-0506 and by ISF grant 1300/08.

†Weizmann Institute of Science. Omer Tamuz is supported by ISF grant 1300/08, and is a recipient of the Google Europe Fellowship in Social Computing. This research is supported in part by this Google Fellowship.

We consider a model describing a group of Bayesian agents that have to make a binary decision. We show that under the dynamics we describe, unanimity is reached with probability one, and give an efficient algorithm for the agents' calculations.

Our model features a finite group of Bayesian agents that have to choose between two possible courses of action. Each initially receives a private and independent signal, which contains some information indicating which action is more likely to be the correct one. The agents participate in rounds of voting, in which each indicates which action it believes is more likely to be correct, and learns the others' opinion thereof. The process continues until unanimity is reached. The Bayesian agents are myopic, so their actions are not strategic, but truthfully reflect the information available to them. They are rational and do not follow heuristics, or rules of thumb, or boundedly-rational courses of action.

As an example, consider a committee that has to decide whether or not to accept a candidate for a position, who a-priori has a chance of one half to be a good hire. Each committee member gets to interview the candidate in private. If the candidate is good - i.e., the correct action is "hire", then each committee member i receives a private signal W_i drawn independently from $N(1, 1)$, the normal distribution with expectation 1 and variance 1. If the candidate is bad (i.e., the correct action is "don't hire"), then W_i is drawn from $N(-1, 1)$.

The committee now commences to vote in rounds. In each round of voting each member casts a public "hire" or "don't hire" vote, depending on which it thinks is more likely to be the correct decision. At each iteration, each member's opinion is based on its private signal, as well as the votes of the other members in the previous rounds.

The agents are Bayesian in the sense that their beliefs are precisely calculated according to Bayes' Law. This is not a straightforward calculation, as they have to take into account that each of the votes of their peers was also likewise calculated. From the description of the process it is not at all clear that there is a succinct description of the decision process taken by the agents, and how can this process be analyzed (indeed - we challenge the reader to try!).

In slightly more general settings the problem seems even more difficult: for example, consider the case that the agents lie on a social network graph, i.e., when each agent only observes the actions of only a subset of the rest. There, no efficient algorithms for the agents' calculations are known, and it is in fact conjectured that none exist (cf., Kanoria and Tamuz [10]). When more than two possible actions are available, then too it seems that the computational problem is significantly more difficult, although perhaps not intractable. An interesting open problem is to suggest an efficient algorithm for the agents' calculation in these (more general) models.

For our model we show that there does exist an efficient algorithm to perform the agents' calculations. We also show that the agents will eventually all cast the same vote, and that the probability that this vote is correct approaches one as the number of agents increases.

1.1 Related work

In 1785 the Marquis de Condorcet proved a founding result [5] in the field of group decision making. The Condorcet Jury Theorem states that given that each member of a jury knows the correct verdict with some probability $p > 1/2$, the probability that the jury reaches a correct decision by a majority vote goes to one as the size of the jury increases. Our "asymptotic learning" result extends this theorem to a more general class of private signals, given that at least two rounds of voting are carried out.

Sebenius and Geanakoplos [19] show in a classical paper that a pair of agents eventually reach agreement on the “state of with world” in a model similar to ours, with finite probability spaces. Likewise, a consequence of the convergence proof given by Gale and Kariv [7]¹ as well as Rosenberg, Solan and Vieille [18] and Mueller-Frank [17] (three models which are generalizations of ours) is that if a pair of agents’ actions converge, it is to the same action, unless the agents are indifferent at the limit. However, none of these results imply that the agents’ actions do in fact converge, or that the agents reach agreement.

We provide a basic proof of a stronger result, namely that all the agents’ actions converge, and in particular to the same action. We further show that each round of voting increases the probability that a given agent votes for the better alternative, and that this probability goes to one at the second round of voting, as the number of participants goes to infinity. Finally, our most significant improvement over the work of Gale and Kariv is that we provide the participants with an efficient algorithm to calculate their beliefs.

The question of complexity of agreement was discussed from a somewhat different perspective in an interesting paper by Aaronson [1]. Aaronson discusses the complexity of agreement in a revealed beliefs model (where agents reveal their beliefs rather than just observe actions), but with general correlated signals. He provides polynomial bounds for approximate convergence (in the sense that the beliefs are close at a specific time t , but with no guarantee on how close they are at later times) and designs even more efficient algorithms for achieving agreement. Of course, in the context of the current paper, if the agents reveal their beliefs, then consensus is achieved in one round. While our paper is much more restricted in terms of the graph (the complete graph) and the signal structure (conditional independence) it considers a natural action dynamics, as opposed to Aaronson’s belief dynamics or artificial dynamics. Our results are also stated in terms of convergence of beliefs and not just closeness of beliefs at a certain point in time.

In a subsequent work to this paper, together with Allan Sly [15], we show that for very general voting models asymptotic learning holds in the sense that as the number of voters goes to infinity the probability of convergence to the correct outcome goes to one. However, the results of this article are not known to extend to the general models studied in [15]. There, no efficient update algorithms are known, no rates of convergence are known, and it is not known whether agents always reach consensus. We answer all of these question for the model we study.

Note that the results of [15] do imply that as the number of agents goes to infinity, the probability of non-convergence (or convergence to the wrong state) goes to zero. However, as in the work of Gale and Kariv, for finite graphs it remains possible that with positive probability both actions are taken infinitely often, with beliefs converging to values implying indifference between the actions. We show that this is not the case in our model.

Studies in committee mechanism design (cf. [12, 9]) strive to construct mechanisms for eliciting information out of committee members to arrive at optimal results. We, by and large, do not take this path but consider a “natural” setting which was not specifically constructed to achieve any such goal.

One could indeed raise the objection that the process could be made simpler if the committee members were to tell each other their private signals, in which case the optimal answer would be arrived at immediately. However, a common assumption in the study of Bayesian agents (cf. [3, 4, 21, 7]) is that “actions speak louder than words”, so that agents learn from each other’s actions rather than revealing to each other all their information. The latter option may not be feasible,

¹See a comment on this paper in [16].

as the said information may consist of experiences and impressions that could take too long to explain, may be difficult to articulate, or may not be even consciously known. In our case the agents' actions are the casting of votes.

Our work is more closely related to models of herd behavior (cf. [3, 4, 21]). These feature a group of agents with a state of the world and corresponding private signals, much like ours. There too agents don't observe each others' private signals but only actions. However, there the agents are exogenously ordered, and each takes a single action after seeing - and learning from - the actions of its predecessors.

1.2 Model

Our model features a finite set of agents $[n] = \{1, 2, \dots, n\}$ that have to make a binary decision regarding an unknown state of the world $S \in \{0, 1\}$. Each is initially given a private signal W_i , distributed μ_0 if $S = 0$ and μ_1 if $S = 1$, and independent of the other signals, conditioned on S .

Definition 1.1. *Let μ_0 and μ_1 be measures on a σ -algebra (Ω, \mathcal{O}) satisfying the following conditions:*

1. μ_0 and μ_1 are mutually absolutely continuous, so that, by the Radon-Nikodym theorem, the Radon-Nikodym derivative $\frac{d\mu_1}{d\mu_0}(\omega)$ exists and is non-zero for all $\omega \in \Omega$.
2. Let W be distributed $\frac{1}{2}\mu_0 + \frac{1}{2}\mu_1$, and let $X = \log \frac{d\mu_1}{d\mu_0}(W)$. Then the distribution of X is non-atomic.

Note that (2) implies that $\mu_0 \neq \mu_1$, since otherwise $X = 0$ a.s. and thus its distribution is atomic.

Definition 1.2. *Let μ_0 and μ_1 be measures on a σ -algebra (Ω, \mathcal{O}) satisfying the conditions of definition 1.1. Let δ_0, δ_1 denote the measures on $\{0, 1\}$ that satisfy $\delta_0(0) = \delta_1(1) = 1$ and $\delta_0(1) = \delta_1(0) = 0$.*

Let \mathbb{P} be the probability measure over the space $\{0, 1\} \times \Omega^n$ given by

$$\mathbb{P} = \frac{1}{2}\delta_0 \otimes \mu_0^{\otimes n} + \frac{1}{2}\delta_1 \otimes \mu_1^{\otimes n}. \quad (1)$$

Let S , taking values in $\{0, 1\}$, and (W_1, \dots, W_n) , taking values in Ω^n , be random variables with joint distribution \mathbb{P} :

$$(S, W_1, \dots, W_n) \sim \mathbb{P}.$$

We call S the state of the world and call W_i agent i 's private signal.

Equivalently, S is picked uniformly from $\{0, 1\}$, and conditioned on S , the agents' private signals W_i are distributed i.i.d. μ_S . It is important to note that conditioned on S the private signals W_i are independent. In much of what follows it is not necessary to assume that they are identical, but we make this assumption to simplify notation and conform to a widely studied economic model.

The agents participate in a process of voting rounds. In each round t each agent i casts a public vote $V_i(t) \in \{0, 1\}$, depending which of the two is more likely to be the state of the world, conditioned on the information available to i ; this includes W_i as well as the votes of the other agents in the previous rounds.

Definition 1.3. For $t \in \{1, 2, \dots\}$ and agent $i \in [n]$, let $V_i(t)$, the vote of agent i at time t , be defined by

$$V_i(t) = \begin{cases} 1 & \text{if } \mathbb{P}[S = 1 | W_i, \bar{V}^{t-1}] > 1/2 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where $\bar{V}^t = \{V_j(t') : j \in [n], t' \leq t\}$ denotes the votes of all agents up to time t .

Alternatively, one could define

$$V_i(t) = \operatorname{argmax}_{s \in \{0,1\}} \mathbb{P}[S = s | W_i, \bar{V}^{t-1}], \quad (3)$$

with a ‘‘tie breaking law’’ specifying that when the conditional probability on the r.h.s. is half then the vote is 0. We note that it is easy to see that the assumption that the distribution of the Radon-Nikodym derivatives $\frac{d\mu_1}{d\mu_0}(W_i)$ is non-atomic (definition 1.1) implies that the probability of ever encountering a tie is 0 a.s. and therefore the details of the tie breaking rule a.s. do not affect the behavior of the process or our results.

1.3 Results

For the model defined in definitions 1.1, 1.2 and 1.3 we prove the following theorems.

- **Unanimity:** A unanimous decision is always reached. That is, with probability one all agents vote identically at some round, and the process essentially ends.

Theorem 1.4. *With probability one there exists a time T_u and a vote $V \in \{0, 1\}$ such that for all $t \geq T_u$ and agents i it holds that $V_i(t) = V$.*

- **Monotonicity:** The probability that an agent votes correctly is non-decreasing with the progression of rounds.

Theorem 1.5. *For all agents i and times $t > 1$, it holds that*

$$\mathbb{P}[V_i(t) = S] \geq \mathbb{P}[V_i(t-1) = S].$$

- **Asymptotic Learning:** The probability of reaching a correct decision at the end of the process approaches one as the number of agents increases. In fact, this already holds by the second round of voting.

Theorem 1.6. *Fix μ_0 and μ_1 , and let n be the number of agents. Then there exist constants $C = C(\mu_0, \mu_1)$ and $n_0 = n_0(\mu_0, \mu_1)$ such that*

$$\mathbb{P}[\forall i : V_i(2) = S] > 1 - e^{-Cn}$$

for all $n > n_0$.

- **Tractability:** In order to discuss tractability we must assume that certain calculations related to the distributions μ_1 and μ_0 take constant time, or alternatively that the algorithm has access to an oracle which performs them in constant time. Specifically, we define below (definition 2.1) the log-likelihood ratio $X = d\mu_1/d\mu_0$ and its conditional distributions ν_0 and ν_1 , and assume that their cumulative distribution functions can be calculated in constant time. Then we show that the agents’ computations are tractable:

Theorem 1.7. Fix μ_0 and μ_1 , and let n be the number of agents. Assume that X , as well as the cumulative distribution functions of ν_0 and ν_1 , can be calculated in constant time. Then there exists an algorithm with running time $O(nt)$, which, given i 's private signal W_i and the votes $\bar{V}^{t-1} = \{V_j(t') : j \in [n], t' < t\}$, calculates $V_i(t)$, agent i 's vote at time t .

We in fact provide a simple algorithm that performs this calculation.

1.4 Comparison to majority voting

Apart from being computationally easier, majority voting is inferior in every one of the above senses. In particular, it doesn't aggregate information as well as repeated voting until consensus, and may not have the **asymptotic learning** property. Consider the following example: A committee has to decide whether or not to accept a candidate for a position. Each member of the committee interviews the candidate and forms an opinion. Now, assume that a good candidate will make a favorable impression nine times out of ten, whereas a bad candidate will make a favorable impression six times out of ten (being good enough to have made it to the interview stage). In this case, with overwhelming probability (i.e., with probability that tends to one as the size of the committee increases), when the candidate is bad, about sixty percent of the committee members will still have a good impression, and consequently a decision by majority will come to the wrong decision, namely that the candidate is good.

This flaw is rectified by a second vote: after seeing the results of the first round of voting, the committee members will realize that too few of them had a good impression, and will vote against the bad candidate in the second round. Indeed, we prove below that asymptotic learning is always achieved after two voting rounds. This suggests that in situations where voting until convergence to consensus is impractical, it may be still be beneficial to have more than one round of voting. Note that there exist other mechanisms that assure efficient aggregation of information. For example, Gerardi and Yariv [8] show that adding a "cheap talk" deliberation phase before a strategic majority vote can also lead to efficient aggregation².

Another characteristic advantage of consensus voting is that the strengths of the participants' convictions counts. Consider a situation in which each agent's private signal is, with high probability, independent of the state of the world, but with some probability provides very convincing evidence. While a single agent possessing the said "smoking gun" would have little impact in a majority vote, his or her insistence in subsequent rounds would convey the weight of the evidence to the rest of the group.

1.5 Asymptotic learning vs. optimal aggregation of information

A stronger notion than asymptotic learning is that of "optimal aggregation of information". This would describe the case that the vote that the agents eventually converge to is equal to the vote that would be cast by a social planner who has access to all the agents' private signals, i.e., $\operatorname{argmax}_{s \in \{0,1\}} \mathbb{P}[S = s | W_1, \dots, W_n]$. We state here without proof that this is not the case in our model.

²Gerardi and Yariv consider strategic agents and show that adding a deliberation phase can lead to equilibria in which information is efficiently aggregated. Essentially, the agents reveal their private signals and then all vote identically. Their work, by its nature, does not consider computational issues.

This question is related to that of monotonicity. Let w_1, \dots, w_n be such that when $W_i = w_i$ then the agents all converge to 1, and let w'_1, \dots, w'_n be such that $\mathbb{P}[S = 1|W_i = w'_i] \geq \mathbb{P}[S = 1|W_i = w_i]$. Is it necessarily the case that setting $W_i = w'_i$ would also result in agreement on 1? Perhaps surprisingly, it is possible to construct examples in which this is not the case. A consequence is that this model does not display optimal aggregation of information, since $\operatorname{argmax}_{s \in \{0,1\}} \mathbb{P}[S = s|W_1, \dots, W_n]$ is monotonic in the sense described above.

1.6 A note on uniform priors

We assume that the agents' prior is uniform, i.e., $\mathbb{P}[S = 1] = \mathbb{P}[S = 0] = \frac{1}{2}$. We make this choice to simplify our notation and make the article easier to read, although our results can be easily extended to biased priors. For this extension an additional requirement is needed: it is not enough that $\mu_0 \neq \mu_1$, since it may be the case that for any value of W_i it holds that $\mathbb{P}[S = 1|W_i] > \frac{1}{2}$. For example, let the prior be such that $\mathbb{P}[S = 1] = 0.9$, and let each private signals W_i equal S with probability 0.51 and equal $1 - S$ with probability 0.49. Then for any value of W_i it will be the case that $V_i(1) = 1$. In this case, although consensus will be reached immediately, there will be no asymptotic learning.

Hence for general priors the requirement is that μ_0, μ_1 be such that $\mathbb{P}[V_i(1) = 1] > 0$ and $\mathbb{P}[V_i(1) = 0] > 0$. Given this, the results above and the ideas of the proof below apply equally.

2 Proofs

Before proving our theorems we make some additional definitions. We start by defining the log-likelihood ratio x and its conditional distributions ν_0 and ν_1 .

Definition 2.1. Let $x : \mathbb{R} \rightarrow \mathbb{R}$ be given by

$$x(\omega) = \log \frac{d\mu_1}{d\mu_0}(\omega). \quad (4)$$

Let ν_0 be the distribution of $x(A)$ when $A \sim \mu_0$ and let ν_1 be the distribution of $x(A)$ when $A \sim \mu_1$.

Note that if $(S, W) \sim \frac{1}{2}\delta_0 \otimes \mu_0 + \frac{1}{2}\delta_1 \otimes \mu_1$, and $X = x(W)$ then

$$x(W) = \log \frac{\mathbb{P}[W|S = 1]}{\mathbb{P}[W|S = 0]},$$

i.e., $x(W)$ is the *log-likelihood ratio* of S given W .

In the proofs that follow we denote

$$X_i = x(W_i)$$

agent i 's private *log-likelihood ratio*. The advantage of log-likelihood ratios is that they are additive for conditionally independent signals.

In our analysis below an event that we often encounter is $a < X_i \leq b$, and hence the following definition will be useful.

Definition 2.2. Let $x : \mathbb{R}^2 \rightarrow \mathbb{R}$ be given by

$$x(a, b) = \log \frac{\mu_1(a < \omega \leq b)}{\mu_0(a < \omega \leq b)}. \quad (5)$$

Note that if $(S, W) \sim \frac{1}{2}\delta_0 \otimes \mu_0 + \frac{1}{2}\delta_1 \otimes \mu_1$, and $X = x(W)$ then

$$x(a, b) = \log \frac{\mathbb{P}[a < X \leq b | S = 1]}{\mathbb{P}[a < X \leq b | S = 0]},$$

i.e., $x(a, b)$ is the log-likelihood ratio of S given that X is between a and b .

The following claim follows by application of Bayes' law to Eq. 5. It follows from this claim that if the cumulative distribution functions of ν_0 and ν_1 can be calculated in constant time then so can $x(\cdot, \cdot)$. In what follows we will need to use the fact that $x(\cdot, \cdot)$ can be efficiently calculated.

Claim 2.3.

$$x(a, b) = \log \frac{\nu_1(X \leq b) - \nu_1(X \leq a)}{\nu_0(X \leq b) - \nu_0(X \leq a)}. \quad (6)$$

We shall also need the following easy claim in some of the proofs below. It can be stated informally as “the log-likelihood ratio of the log-likelihood ratio is the log-likelihood ratio”.

Claim 2.4. Let μ_0, μ_1 be such that $d\mu_1/d\mu_0$ exists and is non-zero for all ω , and let

$$x(\omega) = \log \frac{d\mu_1}{d\mu_0}(\omega).$$

Let ν_0 be the distribution of $x = x(W)$ when $W \sim \mu_0$, and let ν_1 be the distribution of $x = x(W)$ when $W \sim \mu_1$. Then $d\nu_1/d\nu_0$ exists and

$$x = \log \frac{d\nu_1}{d\nu_0}(x). \quad (7)$$

Proof. Since μ_1 and μ_0 are absolutely continuous with respect to each other, it follows from the fact that $x(W)$ is a function of W that ν_1 and ν_0 are also absolutely continuous with respect to each other, and so $d\nu_1/d\nu_0$ exists and is non-zero.

Let $(S, W) \sim \frac{1}{2}\delta_0 \otimes \mu_0 + \frac{1}{2}\delta_1 \otimes \mu_1$, and denote $X = x(W)$. Let

$$M = \mathbb{P}[S = 1 | W] = \mathbb{E}[S | W].$$

Then $X = \log(M/(1 - M))$. By the law of total expectation we have that

$$\mathbb{E}[S | M] = \mathbb{E}[\mathbb{E}[S | M, W] | M].$$

Since M is a function of W then $\mathbb{E}[S | M, W] = \mathbb{E}[S | W] = M$ and it follows that

$$\mathbb{E}[S | M] = \mathbb{E}[M | M] = M.$$

Now from the definition of X it follows that $X = \log(M/(1 - M))$ and therefore there is a one-to-one correspondence between X and M . Hence $\mathbb{E}[S | X] = \mathbb{E}[S | M] = M$. Therefore

$$\log \frac{\mathbb{P}[S = 1 | X]}{\mathbb{P}[S = 0 | X]} = \log \frac{M}{1 - M} = X.$$

But the l.h.s. of this equation is by Bayes' law equal to

$$\log \frac{\mathbb{P}[X|S=1]}{\mathbb{P}[X|S=0]},$$

which is equal to $\frac{d\nu_1}{d\nu_0}(X)$, and thus we have that

$$X = \log \frac{d\nu_1}{d\nu_0}(X).$$

□

2.1 Tractability

The key observation behind our tractability proof is the following. Let i and j be two agents. Since j knows all that i knows except W_i , then even before i votes, agent j can know for which values of W_i agent i would vote 1, and for which it would vote 0. In fact, we show below that there is a bound (that j can calculate) such that if X_i is below that bound then i will vote 0 and otherwise i will vote 1.

Thus at each voting round, j gains either a lower bound or an upper bound on X_i . What we in fact show is that calculating a lower bound $A_i(t)$ and an upper bound $B_i(t)$, on *all* other agents' private likelihood ratios X_i , is almost all that j needs to do to calculate its votes.

The following is the definition of these lower bounds $A_i(t)$ and upper bounds $B_i(t)$.

While Eqs. 8 and 9 may seem mysterious, in what follows we prove that these definitions indeed correspond to the intuition provided above. To somewhat elucidate Eq. 8 (and similarly for Eq. 9) we note that that every time agent i votes 1 the other agents learn a lower bound on X_i ; the best lower bound is the maximum of all these, and therefore we take the maximum over all the times that i voted 1. Each of these bounds depends on the bounds that i learned on the other X_j 's in the previous bounds, and therefore the terms $a_j(\bar{v}^{t'-1})$ and $b_j(\bar{v}^{t'-1})$ appear there.

Definition 2.5. Let $v_i(t) \in \{0, 1\}$ for $i \in [n]$ and $t \in \mathbb{N}$. Similarly to definition 1.3, let $\bar{v}^t = \{v_i(t') : i \in [n], t' \leq t\}$ denote an element of $\{0, 1\}^{nt}$. For $i \in [n]$ and $t \geq 0$, let $a_i : \{0, 1\}^{nt} \rightarrow \mathbb{R}$ and $b_i : \{0, 1\}^{nt} \rightarrow \mathbb{R}$ be the functions recursively defined by

$$a_i(\bar{v}^t) = \max_{t' \leq t \text{ s.t. } v_i(t')=1} \left\{ - \sum_{j \neq i} x \left(a_j(\bar{v}^{t'-1}), b_j(\bar{v}^{t'-1}) \right) \right\}. \quad (8)$$

and

$$b_i(\bar{v}^t) = \min_{t' \leq t \text{ s.t. } v_i(t')=0} \left\{ - \sum_{j \neq i} x \left(a_j(\bar{v}^{t'-1}), b_j(\bar{v}^{t'-1}) \right) \right\}. \quad (9)$$

where the minimum (resp., maximum) over the empty set is taken to be infinity (resp., minus infinity).

Let $A_i(t)$ and $B_i(t)$ be the random variables defined by

$$A_i(t) = a_i(\bar{V}^t)$$

and

$$B_i(t) = b_i(\bar{V}^t).$$

Note that $A_i(t)$ is non-decreasing and $B_i(t)$ is non-increasing, in t .

Since, as we show below, $A_i(t)$ and $B_i(t)$ are lower and upper bounds on X_i at time t , we shall need to often refer to $x(A_i(t), B_i(t))$, and hence denote

$$\hat{X}_i(t) = x(A_i(t), B_i(t)) = x(a_i(\bar{V}^t), b_i(\bar{V}^t)). \quad (10)$$

Recall that $V_i(t)$, agent i 's vote at time t , depends on whether or not $\mathbb{P}[S = 1 | W_i, \bar{V}^{t-1}]$ is greater than half or not:

$$V_i(t) = \begin{cases} 1 & \text{if } \mathbb{P}[S = 1 | W_i, \bar{V}^{t-1}] > 1/2 \\ 0 & \text{otherwise} \end{cases}$$

Let

$$\hat{Y}_i(t) = \log \frac{\mathbb{P}[S = 1 | W_i, \bar{V}^{t-1}]}{\mathbb{P}[S = 0 | W_i, \bar{V}^{t-1}]}. \quad (11)$$

Then

$$V_i(t) = 1 \text{ iff } \hat{Y}_i(t) > 0. \quad (12)$$

Theorem 2.6. *For all $i \in [n]$ and $t > 0$ it holds that*

$$\hat{Y}_i(t) = X_i + \sum_{j \neq i} \hat{X}_j(t-1), \quad (13)$$

Proof. We prove by induction on t . The basis $t = 1$ follows simply from the definitions; since \bar{V}^0 is empty (the agents only start voting at $t = 1$) then $A_i(0) = -\infty$ and $B_i(0) = \infty$ for all $i \in [n]$, and so

$$\hat{X}_i(0) = x(A_i(0), B_i(0)) = x(-\infty, \infty) = 0,$$

by the definition of $x(\cdot, \cdot)$ (Eq. (5)). Another consequence of the fact that \bar{V}^0 is empty is that

$$\hat{Y}_i(1) = \log \frac{\mathbb{P}[S = 1 | W_i]}{\mathbb{P}[S = 0 | W_i]}$$

and so for $t = 1$ the statement of the theorem (Eq. (13)) reduces to

$$\log \frac{\mathbb{P}[S = 1 | W_i]}{\mathbb{P}[S = 0 | W_i]} = X_i,$$

which is precisely the definition of $X_i = x(W_i)$.

Assume the statement holds for all $t' < t$ and all $i \in [n]$. We will show that it holds for t and all i . Since, as we note above, $V_i(t') = 1$ iff $\hat{Y}_i(t') > 0$, then by the inductive assumption we have that

$$V_i(t') = 1 \text{ iff } X_i + \sum_{j \neq i} \hat{X}_j(t' - 1) > 0 \quad (14)$$

or

$$\begin{aligned} V_i(t') &= \mathbf{1} \left(- \sum_{j \neq i} \hat{X}_j(t' - 1) < X_i \right) \\ &= \mathbf{1} \left(- \sum_{j \neq i} x(a_j(\bar{V}^{t'-1}), b_j(\bar{V}^{t'-1})) < X_i \right), \end{aligned}$$

where the second equality follows by substituting the definition of $\hat{X}_i(t')$. Hence $V_i(t')$ is equivalent to either a lower bound (if it equal to 1) or upper bound (if it is equal to 0) on X_i .

Therefore the event $\bar{V}^{t'} = \bar{v}^{t'}$ is equal to the event that

$$X_i > - \sum_{j \neq i} x(a_j(\bar{v}^{t'-1}), b_j(\bar{v}^{t'-1}))$$

for all i and t' such that $v_i(t') = 1$ and

$$X_i \leq - \sum_{j \neq i} x(a_j(\bar{v}^{t'-1}), b_j(\bar{v}^{t'-1}))$$

for all i and t' such that $v_i(t') = 0$. Equivalently, for all $i \in [n]$:

$$\max_{t' \leq t, v_i(t')=1} \left\{ - \sum_{j \neq i} x(a_j(\bar{v}^{t'-1}), b_j(\bar{v}^{t'-1})) \right\} < X_i \leq \min_{t' \leq t, v_i(t')=0} \left\{ - \sum_{j \neq i} x(a_j(\bar{v}^{t'-1}), b_j(\bar{v}^{t'-1})) \right\}$$

Substituting the definitions of a_i (Eq. (8)) and b_i (Eq. (9)), this event is equal to the event

$$a_i(\bar{v}^{t'}) < X_i \leq b_i(\bar{v}^{t'}), \tag{15}$$

for all $i \in [n]$ (note that this means that $A_i(t') < X_i \leq B_i(t')$ for all i and t'). Therefore

$$\mathbb{P} \left[S = s \mid \bar{V}^{t'} = \bar{v}^{t'} \right] = \mathbb{P} \left[S = s \mid a_i(\bar{v}^{t'}) < X_i \leq b_i(\bar{v}^{t'}) \text{ for } i \in [n] \right],$$

and also

$$\mathbb{P} \left[S = s \mid W_i = \omega, \bar{V}^{t'} = \bar{v}^{t'} \right] = \mathbb{P} \left[S = s \mid W_i = \omega, a_j(\bar{v}^{t'}) < X_j \leq b_j(\bar{v}^{t'}) \text{ for } i \neq j \right].$$

Hence

$$\begin{aligned} & \log \frac{\mathbb{P} [S = 1 \mid W_i = \omega, \bar{V}^{t-1} = \bar{v}^{t-1}]}{\mathbb{P} [S = 0 \mid W_i = \omega, \bar{V}^{t-1} = \bar{v}^{t-1}]} \\ &= \log \frac{\mathbb{P} [S = 1 \mid W_i = \omega, a_j(\bar{v}^{t-1}) < X_j \leq b_j(\bar{v}^{t-1}) \text{ for } i \neq j]}{\mathbb{P} [S = 0 \mid W_i = \omega, a_j(\bar{v}^{t-1}) < X_j \leq b_j(\bar{v}^{t-1}) \text{ for } i \neq j]}. \end{aligned}$$

Again invoking Bayes' law, we have that

$$\begin{aligned} & \log \frac{\mathbb{P} [S = 1 \mid W_i = \omega, \bar{V}^{t-1} = \bar{v}^{t-1}]}{\mathbb{P} [S = 0 \mid W_i = \omega, \bar{V}^{t-1} = \bar{v}^{t-1}]} \\ &= \log \left(\frac{\mathbb{P} [W_i = \omega \mid S = 1]}{\mathbb{P} [W_i = \omega \mid S = 0]} \prod_{j \neq i} \frac{\mathbb{P} [a_j(\bar{v}^{t-1}) < X_j \leq b_j(\bar{v}^{t-1}) \mid S = 1]}{\mathbb{P} [a_j(\bar{v}^{t-1}) < X_j \leq b_j(\bar{v}^{t-1}) \mid S = 0]} \right) \end{aligned}$$

since the private signals are independent, conditioned on S . Substituting the definition of $x(\omega)$ (Eq. (4)) and the definition of $x(\cdot, \cdot)$ (Eq. (5)) yields

$$\log \frac{\mathbb{P}[S = 1 | W_i = \omega, \bar{V}^{t-1} = \bar{v}^{t-1}]}{\mathbb{P}[S = 0 | W_i = \omega, \bar{V}^{t-1} = \bar{v}^{t-1}]} = x(\omega) + \sum_{j \neq i} x(a_j(\bar{v}^{t-1}), b_j(\bar{v}^{t-1})). \quad (16)$$

Finally, since

$$\hat{Y}_i(t) = \log \frac{\mathbb{P}[S = 1 | W_i, \bar{V}^{t-1}]}{\mathbb{P}[S = 0 | W_i, \bar{V}^{t-1}]},$$

then by Eq. (16)

$$\hat{Y}_i(t) = x(W_i) + \sum_{j \neq i} x(a_j(\bar{V}^{t-1}), b_j(\bar{V}^{t-1})),$$

and the theorem follows by substituting $X_i = x(W_i)$ and $\hat{X}_j(t-1) = x(a_j(\bar{V}^{t-1}), b_j(\bar{V}^{t-1}))$. \square

We are now ready to prove our main theorem for this subsection.

Theorem 2.7 (Thm. 1.7). *Fix μ_0 and μ_1 , and let n be the number of agents. Assume that X , as well as the cumulative distribution functions of ν_0 and ν_1 , can be calculated in constant time. Then there exists an algorithm with running time $O(nt)$, which, given i 's private signal W_i and the votes $\bar{V}^{t-1} = \{V_j(t') : j \in [n], t' < t\}$, calculates $V_i(t)$, agent i 's vote at time t .*

Proof. By Eq. (12) we have that $V_i(t)$ is a simple function of $\hat{Y}_i(t)$. By Theorem 13 above, $\hat{Y}_i(t)$ can be calculated in $O(n)$ by adding X_i (which we assume can be calculated in constant time given W_i) to the sum over $j \neq i$ of $x(A_j(t), B_j(t))$. By Eq. (6), $x(a, b)$ can be calculated in constant time, assuming the cumulative distribution functions of ν_0 and ν_1 can be calculated in constant time.

We have therefore reduced the problem to that of calculating $A_j(t) = a_j(\bar{V}^t)$ and $B_j(t) = b_j(\bar{V}^t)$. However, the definitions of a_j and b_j (Eqs. (8) and (9)) are in fact simple recursive rules for calculating $a_j(\bar{v}^t)$ and $b_j(\bar{v}^t)$ for all $j \in [n]$, given $a_j^{t-1}(\bar{v}^{t-1})$ and $b_j^{t-1}(\bar{v}^{t-1})$ for all $j \in [n]$: it follows directly from Eqs. (8) and (9) that

$$a_i(\bar{v}^t) = \begin{cases} \max \left\{ a_i(\bar{v}^{t-1}), \sum_{j \neq i} x(a_j(\bar{v}^{t-1}), b_j(\bar{v}^{t-1})) \right\} & \text{if } v_i(t) = 1 \\ a_i(\bar{v}^{t-1}) & \text{otherwise} \end{cases},$$

with an analogous equation for $b_i(\bar{v}^t)$.

Note that the sum $\sum_{j \neq i} x(a_j(\bar{v}^{t-1}), b_j(\bar{v}^{t-1}))$ needn't be calculated from scratch for each i ; one can rather sum over all $j \in [n]$ once and subtract the appropriate term for each i . Hence calculating $a_j(\bar{v}^t)$ and $b_j(\bar{v}^t)$ (for all j) given their predecessors takes $O(n)$, and the entire recursive calculation takes $O(nt)$. \square

2.2 Unanimity

Our model is a special case of that of Gale and Kariv [7]. They show a weak agreement result: namely, that if the votes of two agents converge, and if the agents are not indifferent at the limit $t \rightarrow \infty$, then they converge to the same vote. We prove the strongest possible agreement result: consensus is reached with probability one, i.e., the agents almost always all converge to the same vote.

Before proving the theorem we prove some standard claims. Recall the definition of $x(\cdot, \cdot)$ (Eq. (5)):

$$x(a, b) = \log \frac{\mathbb{P}[S = 1 | a < X \leq b]}{\mathbb{P}[S = 0 | a < X \leq b]}.$$

Claim 2.8. *Let a, b be such that $x(a, b)$ is well defined (i.e., $\mathbb{P}[a < X \leq b | S = 0] > 0$). Then*

$$x(a, b) = \log \mathbb{E} [e^X | a < X \leq b, S = 0]. \quad (17)$$

Proof. By Bayes' law we have that

$$x(a, b) = \log \frac{\mathbb{P}[a < X \leq b | S = 1]}{\mathbb{P}[a < X \leq b | S = 0]},$$

Substituting the conditional distributions of X yields

$$x(a, b) = \log \frac{\int_a^b d\nu_1(x)}{\int_a^b d\nu_0(x)}.$$

By Claim 2.4

$$x = \log \frac{d\nu_1}{d\nu_0}(x), \quad (18)$$

and so we have that

$$x(a, b) = \log \frac{\int_a^b \frac{d\nu_1}{d\nu_0}(x) d\nu_0(x)}{\int_a^b d\nu_0(x)} = \log \frac{\int_a^b e^x d\nu_0(x)}{\int_a^b d\nu_0(x)}.$$

Recalling that ν_0 is the distribution of X conditioned on $S = 0$ we have that

$$x(a, b) = \log \mathbb{E} [e^X | a < X \leq b, S = 0].$$

□

Recall that we assume that the distribution of X is non-atomic (definition 1.1). Hence the following claim is a consequence of Eq. (17) above, by a standard argument that we omit.

Claim 2.9. *$x(a, b)$ is non-decreasing and continuous in a and in b .*

The following claims follows directly from Eq. (17) above.

Claim 2.10. *Let a, b be such that $x(a, b)$ is well defined (i.e., $\mathbb{P}[a < X \leq b | S = 0] > 0$). Then $a < x(a, b) < b$, assuming the distribution of X is non-atomic.*

Proof. By Eq. (17) we have that

$$e^{x(a,b)} = \mathbb{E} \left[e^X \mid e^a < e^X \leq e^b, S = 0 \right],$$

and so $e^a < e^{x(a,b)} \leq e^b$. Since we assume the distribution of X is non-atomic (definition 1.1) then

$$\mathbb{E} \left[e^X \mid e^a < e^X \leq e^b, S = 0 \right] < e^b,$$

and so $e^a < e^{x(a,b)} < e^b$ and the claim follows. \square

We now show a condition for unanimity. We will later prove that unanimity occurs w.p. 1 by showing that this condition eventually applies, w.p. 1.

Lemma 2.11. *If*

$$\sum_i (B_i(t) - A_i(t)) < \left| \sum_i X_i \right| \tag{19}$$

then there exists a V such that $V_i(t') = V$ for all i and $t' > t$. I.e., unanimity is reached at time t .

Proof. We first note that, since $A_i(t)$ is non-decreasing and $B_i(t)$ is non-increasing then if Eq. 19 holds at time t then it also holds at all times $t' > t$.

Now, recall that $\hat{X}_i(t) = x(A_i(t), B_i(t))$. By Claim 2.10 we have that $A_i(t) < \hat{X}_i(t) \leq B_i(t)$. From Eq. (15) it follows that the same holds for X_i too: $A_i(t) < X_i \leq B_i(t)$. Hence $|\hat{X}_i(t) - X_i| \leq B_i(t) - A_i(t)$ and for all $i \in [n]$ we have that

$$\left| \sum_{j \neq i} (\hat{X}_j(t) - X_j) \right| \leq \sum_j (B_j(t) - A_j(t)). \tag{20}$$

Recall that

$$\hat{Y}_i(t+1) = X_i + \sum_{j \neq i} \hat{X}_j(t),$$

and so

$$\hat{Y}_i(t+1) - \sum_j X_j = \sum_{j \neq i} (\hat{X}_j(t) - X_j).$$

Therefore by Eq. (20) we have that

$$\left| \hat{Y}_i(t+1) - \sum_j X_j \right| \leq \sum_j (B_j(t) - A_j(t)).$$

By the theorem hypothesis this implies that

$$\left| \hat{Y}_i(t+1) - \sum_j X_j \right| \leq \left| \sum_j X_j \right|.$$

Hence $\hat{Y}_i(t+1)$ and $\sum_j X_j$ have the same sign. Since $V_i(t+1) = 1$ iff $\hat{Y}_i(t+1) > 0$ (Eq. (12)) then we have shown that at time $t+1$ all agents vote identically. Since if Eq. 19 holds for time t then it also holds for time $t+1$ then we've shown that for all $t' > t$ the agents will agree in every round. It remains to show that they don't all change their opinion, as a group.

Now, if the agents all vote 1 at time $t+1$ then, by the definition of $A_i(t)$ and $B_i(t)$, it holds that $B_i(t+2) = B_i(t+1)$ and $A_i(t+2) \geq A_i(t+1)$. Since by Claim 2.9 $x(a, b)$ is non-decreasing in a , then we have that $\hat{X}_i(t+2) \geq \hat{X}_i(t+1)$ for all i , and so $\hat{Y}_i(t+2) \geq \hat{Y}_i(t+1)$ for all i . Hence the agents will all vote 1 at time $t+2$. The same argument applies when all the agents vote 0 at time $t+1$, and the proof follows by induction on t . \square

We make another definition before proceeding to prove the main theorem of this subsection. Recall the definitions of a_i , b_i , A_i and B_i :

$$a_i(\bar{v}^t) = \max_{t' \leq t \text{ s.t. } v_i(t')=1} \left\{ - \sum_{j \neq i} x \left(a_j(\bar{v}^{t'-1}), b_j(\bar{v}^{t'-1}) \right) \right\}$$

and

$$b_i(\bar{v}^t) = \min_{t' \leq t \text{ s.t. } v_i(t')=0} \left\{ - \sum_{j \neq i} x \left(a_j(\bar{v}^{t'-1}), b_j(\bar{v}^{t'-1}) \right) \right\}.$$

with $A_i(t) = a_i(V^t)$ and $B_i(t) = b_i(V^t)$. As we noted above $A_i(t)$ is non-decreasing in t and $B_i(t)$ is non-increasing in t . Hence they have limits which we denote by $A_i(\infty)$ and $B_i(\infty)$. Furthermore, if as above we denote $\hat{X}_i(t) = x \left(a_j(\bar{V}^{t'}), b_j(\bar{V}^{t'}) \right)$ then

$$A_i(\infty) = \sup_{t' \text{ s.t. } V_i(t')=1} \left\{ - \sum_{j \neq i} \hat{X}_i(t') \right\} \quad (21)$$

and

$$B_i(\infty) = \inf_{t' \text{ s.t. } V_i(t')=0} \left\{ - \sum_{j \neq i} \hat{X}_i(t') \right\} \quad (22)$$

Note that since $\hat{X}_i(t) = x(A_i(t), B_i(t))$, and since $x(a, b)$ is a continuous function of a and b (Claim 2.9) then

$$\lim_{t \rightarrow \infty} \hat{X}_i(t) = x(A_i(\infty), B_i(\infty)). \quad (23)$$

Theorem 2.12 (Thm. 1.4). *With probability one there exists a time T_u and a vote $V \in \{0, 1\}$ such that for all $t \geq T_u$ and agents i it holds that $V_i(t) = V$.*

Proof. Assume by way of contradiction that unanimity is never reached, and so by Lemma 2.11 for all t it holds that $\sum_i B_i(t) - A_i(t) \geq |\sum_i X_i|$. Then, since $B_i(t) - A_i(t)$ is monotonically decreasing, it holds that

$$\lim_{t \rightarrow \infty} \sum_i (B_i(t) - A_i(t)) \geq \left| \sum_i X_i \right|. \quad (24)$$

Let $Z := \lim_{t \rightarrow \infty} \sum_j \hat{X}_j(t)$. We consider separately the events that $Z = 0$, $Z < 0$ and $Z > 0$:

1. $Z = 0$

We assume (definition 1.1) that the distribution of X_i is non-atomic, and so $\sum_i X_i \neq 0$ with probability 1. Hence by Eq. (24) there must be some agent i for which

$$\lim_{t \rightarrow \infty} (B_i(t) - A_i(t)) = B_i(\infty) - A_i(\infty) > 0.$$

Assume w.l.o.g. $V_i(t) = 1$ infinitely many times. Hence, by Eq. (21), we have that

$$\begin{aligned} A_i(\infty) &\geq - \lim_{t \rightarrow \infty} \sum_{j \neq i} \hat{X}_j(t) \\ &= \lim_{t \rightarrow \infty} \hat{X}_i(t) - \lim_{t \rightarrow \infty} \sum_j \hat{X}_j(t). \end{aligned}$$

Since we assume in this case that $Z = \lim_{t \rightarrow \infty} \sum_j \hat{X}_j(t) = 0$ then we have that

$$A_i(\infty) \geq \lim_{t \rightarrow \infty} \hat{X}_i(t).$$

But since $A_i(\infty) < B_i(\infty)$ then by Eq. (23) and Claim 2.10 we have that $A_i(\infty) < \lim_{t \rightarrow \infty} \hat{X}_i(t)$, which is a contradiction.

The intuition here is that when i votes 1 it is revealed that $X_i > \hat{X}_i(t) - \sum_j \hat{X}_j(t)$. Hence if $\sum_j \hat{X}_j(t)$ is very small then $A_i(t)$ approaches $\hat{X}_i(t)$ arbitrarily closely, which is impossible if $A_i(t)$ is to stay well separated from $B_i(t)$.

2. $Z > 0$

Since unanimity is never reached then there must be some i for which $\hat{Y}_i(t) \leq 0$ infinitely many times. Hence by Eq. (22) we have that

$$\begin{aligned} B_i(\infty) &\leq - \lim_{t \rightarrow \infty} \sum_{j \neq i} \hat{X}_j(t) \\ &= \lim_{t \rightarrow \infty} \hat{X}_i(t) - \lim_{t \rightarrow \infty} \sum_j \hat{X}_j(t). \end{aligned}$$

Since by assumption $\lim_{t \rightarrow \infty} \sum_j \hat{X}_j(t) > 0$ then we have that

$$B_i(\infty) < \lim_{t \rightarrow \infty} \hat{X}_i(t).$$

However by Eq. (23) and Claim 2.10 we have that $B_i(\infty) > \lim_{t \rightarrow \infty} \hat{X}_i(t)$, which is a contradiction.

In this case the intuition is that when i votes 0 even though $\sum_j \hat{X}_j(t)$ is positive, then B_i decreases by at least $\sum_j \hat{X}_j(t)$, which cannot continue indefinitely when $\lim_{t \rightarrow \infty} \sum_j \hat{X}_j(t) > 0$.

3. $Z < 0$

The argument here is identical to that of the previous case.

□

2.3 Monotonicity

Since the agents base their decisions on a growing information base, their decisions become more and more likely to be correct. We prove this formally below, using a standard argument.

Theorem 2.13 (Thm. 1.5). *For all agents i and times $t > 1$, it holds that*

$$\mathbb{P}[V_i(t) = S] \geq \mathbb{P}[V_i(t-1) = S].$$

Proof. As noted in Eq. (3), $V_i(t)$ is the choice in $\{0, 1\}$ that maximizes the probability of matching the state of the world, given W_i and \bar{V}^{t-1} . Let f be an arbitrary function of W_i and \bar{V}^{t-1} . Then:

$$\mathbb{P}[V_i(t) = S | W_i, \bar{V}^{t-1}] \geq \mathbb{P}[f(W_i, \bar{V}^{t-1}) = S | W_i, \bar{V}^{t-1}].$$

Since $V_i(t-1)$ is also a function of W_i and \bar{V}^{t-1} then we can substitute $V_i(t-1)$ for $f(W_i, \bar{V}^{t-1})$ in the equation above, and the theorem follows. \square

Note that $\mathbb{P}[V_i(t) = S]$ is strictly larger than $\mathbb{P}[V_i(t-1) = S]$ whenever $\mathbb{P}[V_i(t) \neq V_i(t-1)]$ is positive, i.e. when the decision may change.

2.4 Asymptotic Learning

We show that with high probability after observing the first round of voting all voters know the correct state of the world, and a unanimous and correct decision is reached at the second round of voting. Note that by the monotonicity theorem (1.5), this means that the same holds for all rounds after the second round.

Before proving this theorem we will prove the following claim.

Claim 2.14.

$$\mathbb{P}[X < -C | S = 1] < e^{-C}.$$

Proof. Recall that by Claim 2.4

$$x = \frac{d\nu_1}{d\nu_0}(x)$$

and so

$$\int_{-\infty}^{\infty} d\nu_0(X) = \int_{-\infty}^{\infty} e^{-X} d\nu_1(X).$$

But ν_0 is a probability measure, and so $\int_{-\infty}^{\infty} d\nu_0(X) = 1$. Hence

$$\mathbb{E}[e^{-X} | S = 1] = \int_{-\infty}^{\infty} e^{-X} d\nu_1(X) = 1.$$

Therefore by the Markov bound

$$\mathbb{P}[X < -C | S = 1] = \mathbb{P}[e^{-X} > e^C | S = 1] < e^{-C}.$$

\square

We are now ready to prove the main theorem of this subsection.

Theorem 2.15 (Thm. 1.6). *Fix μ_0 and μ_1 , and let n be the number of agents. Then there exist constants $C = C(\mu_0, \mu_1)$ and $n_0 = n_0(\mu_0, \mu_1)$ such that*

$$\mathbb{P}[\forall i : V_i(2) = S] > 1 - e^{-Cn}$$

for all $n > n_0$.

Proof. We shall show that there exists a constant $C = C(\mu_0, \mu_1)$ such that

$$\mathbb{P}[\forall i : V_i(2) = 1 | S = 1] > 1 - e^{-Cn}$$

for all n large enough. Since the same argument can be used to show an analogous statement for $S = 0$ then this will prove the theorem.

Recall that $V_i(t)$ is the indicator of the event $\hat{Y}_i(t) > 0$, where

$$\hat{Y}_i(t) = \log \frac{\mathbb{P}[S = 1 | W_i, \bar{V}^{t-1}]}{\mathbb{P}[S = 0 | W_i, \bar{V}^{t-1}]}.$$

Invoking Bayes' law and the conditional independence of the private signals we get that

$$\hat{Y}_i(2) = \log \left(\frac{\mathbb{P}[W_i | S = 1]}{\mathbb{P}[W_i | S = 0]} \prod_{j \neq i} \frac{\mathbb{P}[V_j(1) | S = 1]}{\mathbb{P}[V_j(1) | S = 0]} \right)$$

Denote by N_i the number of agents other than i who vote 1 in the first round:

$$N_i = |\{j \text{ s.t. } V_j(1) = 1, j \neq i\}|.$$

Then since $X_i = \log \frac{\mathbb{P}[W_i | S=1]}{\mathbb{P}[W_i | S=0]}$ then we can write

$$\hat{Y}_i(2) = X_i + N_i \log \frac{\mathbb{P}[X > 0 | S = 1]}{\mathbb{P}[X > 0 | S = 0]} + (n - 1 - N_i) \log \frac{\mathbb{P}[X \leq 0 | S = 1]}{\mathbb{P}[X \leq 0 | S = 0]}.$$

Denote

$$\alpha_1 = \mathbb{P}[X > 0 | S = 1] \quad \text{and} \quad \alpha_0 = \mathbb{P}[X > 0 | S = 0],$$

and note that $x(0, \infty) > 0$ by Claim 2.10, and so since $x(0, \infty) = \log \alpha_1 / \alpha_0$ then $\alpha_1 \neq \alpha_0$. We can now write

$$\hat{Y}_i(2) = X_i + N_i \log \frac{\alpha_1}{\alpha_0} + (n - 1 - N_i) \log \frac{1 - \alpha_1}{1 - \alpha_0}.$$

Denote

$$W_i = N_i \log \frac{\alpha_1}{\alpha_0} + (n - 1 - N_i) \log \frac{1 - \alpha_1}{1 - \alpha_0}$$

so that $\hat{Y}_i(2) = X_i + W_i$. Since $\mathbb{E}[N_i|S = 1] = (n-1)\alpha_1$ then the conditioning W_i on $S = 1$ we get that

$$\mathbb{E}[W_i|S = 1] = (n-1)\alpha_1 \log \frac{\alpha_1}{\alpha_0} + (n-1)(1-\alpha_1) \log \frac{1-\alpha_1}{1-\alpha_0}.$$

If we denote $(n-1)D = \mathbb{E}[W_i|S = 1]$ then D is the Kullback-Leibler divergence [11] of two Bernoulli distributions with expectations $\alpha_1 \neq \alpha_0$. Hence $D > 0$.

Now, conditioned on S the private signals are independent, and hence so are the votes at round 1, since $V_i(1)$ depends on W_i only. Therefore W_i , conditioned on $S = 1$, is the sum of $n-1$ bounded independent random variables. Therefore by the Hoeffding bound there exists a constant C_1 such that

$$\mathbb{P}[W_i \leq \frac{1}{2}(n-1)D | S = 1] \leq e^{-C_1(n-1)}. \quad (25)$$

By Claim 2.14 we have that $\mathbb{P}[X_i \leq -\frac{1}{2}(n-1)D | S = 1] < e^{-\frac{1}{2}D(n-1)}$, which, together with Eq. (25) and the union bound yields

$$\mathbb{P}[V_i(2) = 0 | S = 1] < e^{-\frac{1}{2}D(n-1)} + e^{-C_1(n-1)}.$$

Therefore, again using the union bound it follows that

$$\mathbb{P}[\exists i : V_i(2) = 0 | S = 1] < n(e^{-\frac{1}{2}D(n-1)} + e^{-C_1(n-1)}).$$

Finally, it follows that for n large enough there exists a constant C such that

$$\mathbb{P}[\forall i : V_i(2) = 1 | S = 1] > 1 - e^{-Cn}.$$

□

References

- [1] S. Aaronson. The complexity of agreement. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 634–643. ACM, 2005.
- [2] R. J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- [3] A. V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- [4] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, pages 992–1026, 1992.
- [5] J.-A.-N. Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. De l'Imprimerie Royale, 1785.
- [6] European Union. Treaty of Lisbon, December 2007. Article 16(4). Available online at http://europa.eu/lisbon_treaty/index_en.htm.

- [7] D. Gale and S. Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, November 2003.
- [8] D. Gerardi and L. Yariv. Deliberative voting. *Journal of Economic Theory*, 134(1):317–338, 2007.
- [9] J. Glazer and A. Rubinstein. Motives and implementation: On the design of mechanisms to elicit opinions. *Journal of Economic Theory*, 79(2):157–173, 1998.
- [10] Y. Kanoria and O. Tamuz. Tractable bayesian social learning on trees. To appear in the proceedings of the *IEEE 2012 International Symposium on Information Theory*, 2012.
- [11] S. Kullback and R. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [12] J. Laslier and J. Weibull. Committee decisions: optimality and equilibrium. *Working Paper Series in Economics and Finance*, 2008.
- [13] R. D. McKelvey and T. Page. Common knowledge, consensus, and aggregate information. *Econometrica*, 54(1):109–127, 1986.
- [14] F. P. Miller, A. F. Vandome, and J. McBrewster. *Consensus Decision Making*. Alphascript Publishing, 2009.
- [15] E. Mossel, A. Sly, and O. Tamuz. Asymptotic learning on bayesian social networks. To appear in *Probability Theory and Related Fields*, 2012.
- [16] M. Mueller-Frank. A comment on ‘bayesian learning in social networks’. Available at SSRN: <http://ssrn.com/abstract=1690943> or <http://dx.doi.org/10.2139/ssrn.1690943>, 2010.
- [17] M. Mueller-Frank. A general framework for rational learning in social networks. *Theoretical Economics*, 8(1), 2013.
- [18] D. Rosenberg, E. Solan, and N. Vieille. Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979–994, 2009.
- [19] J. Sebenius and J. Geanakoplos. Don’t bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424–426, 1983.
- [20] H. Simon. *Reason in Human Affairs*. Stanford University Press, 1982.
- [21] L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.