

On Agreement and Learning*

Elchanan Mossel[†], Allan Sly[‡] and Omer Tamuz[§]

July 26, 2012

Abstract

We consider a group of Bayesian agents who each possess an independent private signal about an unknown state of the world. We study the question of *efficient learning*: in which games is private information efficiently disseminated among the agents? In particular, we explore the notion of *asymptotic learning*, which is said to occur when agents learn the state of the world with probability that approaches one as the number of agents tends to infinity.

We show that under general conditions asymptotic learning follows from agreement on posterior actions or posterior beliefs, regardless of the information exchange dynamics.

Keywords: Bayesian learning, rational expectations, aggregation of information.

1 Introduction

1.1 Overview

Following Aumann’s Agreement Theorem [1], the study of the exchange of information between Bayesian agents has resulted in broad theoretical insight into the phenomenon of agreement and the dynamics that lead to it. However, while questions regarding convergence to common posterior beliefs or common posterior actions are now well understood, the correctness or optimality of these beliefs/actions are understood only in special cases. We present general results showing “asymptotic learning” for agents who are fully Bayesian, and iteratively act or communicate until they converge to the same posterior belief/action.

We consider a group of Bayesian agents who have to make a choice between two alternative actions. The agents are initially given informative, independent and identically distributed private signals. They then participate in a general game or process, during which they learn from their peers. We are not interested in the details of the dynamics of the game, but in what can be said under the condition that at its conclusion the agents reach agreement. We consider two cases: in the first they agree on posterior beliefs regarding the optimal action, i.e. the posterior probability

*We would like to thank Robert J. Aumann for his valuable comments on earlier drafts, which allowed us to extend the statement of some of our results to the common knowledge framework.

[†]Weizmann Institute and U.C. Berkeley. E-mail: mossel@stat.berkeley.edu. Supported by NSF Career Award (DMS 054829), by ONR award N000141110140 and by ISF grant 1300/08.

[‡]Microsoft Research, Redmond, WA.

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

that the first action is optimal. In the second they merely agree on which action is more likely to be optimal.

For large groups of agents the combined private signals contain enough information to identify the optimal action with probability approaching one. It is natural to ask whether the agents learn, i.e., whether this information reaches the agents, or whether information cascades may emerge, with arbitrarily large groups of agents who make the wrong decision with probability bounded away from zero. We show that “asymptotic learning” occurs for a wide variety of models, so that as the size of the group grows the agents learn the correct action with probability that approaches one.

Regardless of the game the agents play, we show that if a group of agents participates in any game or process which always results in agreement on posterior beliefs, then information is aggregated efficiently and asymptotic learning holds. We further show that when private beliefs are unbounded then agreement on actions is a sufficient condition for asymptotic learning.

Many studies describe Bayesian games or dynamics which always result in agreement, and to which, therefore, these results apply. For example, consider a group of agents who are members of a social network and iteratively reveal their true posterior beliefs to their neighbors. This dynamics almost surely leads to common knowledge and therefore agreement on posterior beliefs, as shown by Parikh and Krasucki [17]. Rosenberg, Solan and Vieille [18] show that this still holds when the agents are strategic with some discount factor, in any equilibrium. Simpler mechanisms also exist: For example, McKelvey and Page [12] show that agreement on posterior beliefs can be achieved by repeated public announcements of certain statistics of the beliefs. Our results apply and establish learning in all these cases.

The synthesis of this work with that of Rosenberg, Solan and Vieille [18] may be of particular interest. They describe (among a number of models they consider) an iterative game involving strategic Bayesian agents on a social network who are trying to estimate a binary state of the world from initial private signals. They show that when the network is strongly connected then the result is convergence to a common posterior belief, in any equilibrium. Our theorem applies to the game they describe, for the case of a binary state of the world and conditionally independent private signals.

Feigenbaum et al. [8] and Ostrovsky [16] show efficient aggregation of information in some particular market games. Our results imply that when private signals are conditionally independent then in *any* Bayesian market game in which there exist equilibrium prices - i.e., agents reach agreement - these prices exactly reflect the information available to society.

1.2 General framework

We consider three models which share a common basis: There is a finite number of agents, with $S \in \{0, 1\}$ a binary state of the world which the agents are interested in knowing. Each of the two possible states of the world occurs with probability one half.

The agents are initially provided with signals which are informative with respect to S and independent, conditioned on S : There are two distributions, $\mu_0 \neq \mu_1$, such that conditioned on S the private signals are independent and distributed as μ_S . In Section 1.3.1 we slightly depart from this framework by weakening the assumption on the private signals from being independent to having zero correlation.

In all models the agents are assumed to be Bayesian and participate in some general game or process involving the exchange of information. We are not interested in the actual dynamics, but

only in the informational state at the conclusion of the process.

We define the action of an agent to be “1” if, at the conclusion of the process it considers the probability of the event “ $S = 1$ ” to be higher than one half, and “0” otherwise. Asymptotic learning occurs when the probability that all the agents’ actions are equal to the true state of the world S tends to one as the number of agents tends to infinity, for a fixed choice of private signal distributions μ_1 and μ_0 .

1.3 Common posterior beliefs

Aumann’s seminal Agreement Theorem [1] considers two agents who, following a general game or process, have posterior beliefs regarding some event which are common knowledge. Regardless of how this state was reached, the theorem states that their posterior beliefs must be equal.

Subsequent results by Geanakoplos and Polemarchakis [10] showed that two rational agents, who repeatedly share their beliefs regarding the probability of some event, must converge to identical opinions and cannot “agree to disagree”. This result was later extended to apply to a group of agents connected by a social network by Parikh and Krasucki [17], and furthermore to strategic agents by Rosenberg, Solan and Vieille [18]. The spirit of these findings is that strategic Bayesian agents who communicate their beliefs eventually reach agreement, under a very wide spectrum of dynamics.

However, only a few results consider the value of the common posterior belief, and these apply to particular games and signal distributions. In the case of independent Gaussian signals and myopic agents on a social network, DeMarzo et al. [6] showed that information is optimally aggregated, in the sense that the agents converge to the beliefs they would have if they had access to all the available information. Furthermore, the Bayesian updates in this case are simple and converge rapidly, as shown by DeMarzo et al. [6] and Mossel and Tamuz [15]. Other results showing optimal aggregation are those of Feigenbaum et al. [8] and Ostrovsky [16], for very general signal distributions, but specific models of dynamic market games.

The first question we study in this paper is the following: How effectively do agents who reach agreement on beliefs aggregate their private information? Note that we follow Aumann [1] in considering only the informational state at the conclusion of the game and disregarding the details of the process.

The common posterior belief depends on all the private signals - the different pieces of information that were initially available to the individuals - as well as the dynamics of the interaction. For large enough groups, the information contained in all the private signals suffices to determine the correct action, except with exponentially low probability (in the number of agents). Is it possible that this information is not disseminated to the agents?

We show that it is impossible to reach agreement on posterior beliefs without learning a great deal in the process: in fact, in this case information is optimally aggregated, in the sense that the common posterior belief is the same as it would be if all agents observed all the private signals. In particular, this means that “asymptotic learning” occurs: as the size of the group tends to infinity the probability that the agents’ common action will equal the state of the world tends (exponentially fast) to one, irrespective of the details of the game.

We note that while the assumption of independence is critical for this result, it is not imperative that all the signals be identical; our proof of this theorem holds, as is, for the case that signals are conditionally independent but not identical.

1.3.1 Uncorrelated private signals

The assumption of conditional independence of private signals may be too strong in certain settings. However, when it is completely discarded then asymptotic learning may no longer be a consequence of agreement, and in fact it may be the case that agents agree while learning nothing at all from each other (see example 2.7 below). Hence it may be of interest to explore how this assumption can be weakened while preserving asymptotic learning. Specifically, we consider the case that private signals are conditionally *uncorrelated*¹. For example, signals that are pairwise conditionally independent (but perhaps otherwise dependent) are uncorrelated.

We show that in this case too, agreement implies asymptotic learning. In particular, we show that if the agents agree on a posterior belief - and therefore on an action - then the action they choose is optimal with probability that is inversely proportional to the number of agents n . We show in Example 2.8 that this result is tight.

1.4 Common posterior actions

We here again consider a binary state of the world, conditionally i.i.d. private signals, and a general information exchange process, as in Section 1.3. However, the agents in this case are not guaranteed to learn each others' posterior beliefs (i.e., posterior probabilities that $S = 1$), but rather their *posterior actions*: recall that the action is “0” if the posterior probability of $S = 1$ is less than half and “1” if this probability is greater than half. If this probability equals half then both actions are optimal, and we don't make any assumptions about how the agents break the tie. We therefore, in fact, speak of common knowledge of an agent's “optimal action set” L , which may equal $\{0\}$, $\{1\}$ or $\{0, 1\}$.

Common knowledge of actions is in general an insufficient condition for asymptotic learning. Indeed, there are examples in which common knowledge of actions is achieved without asymptotic learning, typically involving a small group of agents whose opinions dominate those of the remaining agents. However, these examples do not include “highly opinionated” agents whose initial private beliefs may come arbitrarily close to 0 or 1. Indeed, we show that under the condition of “unbounded private beliefs”² asymptotic learning occurs whenever there is common knowledge of actions, regardless of the dynamics. That is, the agents' common optimal action set L equals $\{S\}$ with probability that goes to one as the number of agents goes to infinity. This result shares the flavor of other results (cf. Smith and Sørensen [20]) that rule out information cascades when private beliefs are unbounded.

This result applies to many models in which, through varied dynamics, agents reach common optimal action sets, for the case of the above described structure of private signals. These include models studied by Ménager [13], Sebenius and Geanakoplos [19] and one of the games studied by Rosenberg, Solan and Vieille [18]. In that game, under certain conditions, strategic agents converge to the same action in any equilibrium. Thus this result shows that learning is efficient in any equilibrium of their game, provided that private signals are conditionally independent.

Related questions have been studied in particular games. One line of work considers the case

¹The precise condition for is that the log-likelihood ratio of private beliefs has zero covariance, conditioned on S . We define this more formally in Definition 2.6 below.

²We define unbounded private beliefs in Definition 2.4 below. Essentially, private beliefs are unbounded when private signals inspire beliefs that may be arbitrarily close to both zero and one. Note that in general there is no requirement that beliefs are ever actually equal to zero or one.

where each Bayesian agent acts only once, after observing the actions of a subset of the other agents. In this case “information cascades” may emerge (cf. Banerjee [3], Bikhchandani, Hirshleifer and Welch [5]): There exist arbitrarily large groups of agents who take the wrong action with probability bounded away from zero. Here too a condition for asymptotic learning (i.e., no information cascade) is unbounded private beliefs, as Smith and Sørensen [20] show. Note that while their results are similar in flavor (unbounded beliefs imply asymptotic learning), they study a particular game in which, as they show, agreement on actions is reached, while we consider *any* such game.

A second line of work is the one explored by Ellison and Fudenberg [7] and Bala and Goyal [2], who each study a particular non-Bayesian game involving social networks. The latter are also interested in asymptotic learning (or “complete learning”, in their terms), and show some results of asymptotic learning and some results of non-learning, depending on various parameters of their model. Since their models are not Bayesian they are very different in nature, making it difficult to draw comparisons between our results and theirs.

A paper by Minehart and Scotchmer [14] studies the same structure of the state of the world and private signals. They consider a solution concept they call *Rational Expectations Equilibrium*, and show that when it exists then agents agree on actions, and that when private signals are unbounded then the agreed on action is correct with probability that approaches one as the number of agents tend to infinity. Despite the superficial similarity, however, their model differs from ours in several crucial aspects, and includes what is a strictly weaker result.

In their model posterior probabilities are factored in terms of the actions of the different players (their Eq. 1). This would correspond to a situation where agents in equilibrium have no information regarding the others’ private signals, except what is inferred from their posterior actions. Although this setting is amenable to analysis, for a given distribution of private signals such an equilibrium may not always exist, and indeed they note that typically this is the case (their Proposition 2). In particular, their model does not describe general games or processes in which agents learn more than the “final” posterior action. These include the games described by Rosenberg, Solan and Vieille [18] and the model of Ménager [13], to both of which our results do apply.

2 Formal definitions, results and examples

2.1 Definitions

Definition 2.1. Agents: *We consider a finite set of n agents denoted by V .*

Definition 2.2. State of the worlds and signals: *We assume the state of the world is binary $S \in \{0, 1\}$ and that the distribution of S is uniform: $\mathbb{P}[S = 0] = \mathbb{P}[S = 1] = 1/2$.*

Given a set of agents V as in Definition 2.1 and a state of the world $S \in \{0, 1\}$, we define the private signals of the agents $\{W_u : u \in V\} \in \Omega^V$ as follows. When $S = 1$ (respectively, $S = 0$) then the signals are distributed i.i.d. according to μ_1 (respectively μ_0), where μ_1, μ_0 are two distributions over the same probability space (Ω, \mathcal{O}) . We assume throughout that the private signals are informative, i.e., $\mu_1 \neq \mu_0$.

We assume that the Radon-Nikodym derivative $d\mu_1/d\mu_0$ exists for all $\omega \in \Omega$. We denote by z the log-likelihood ratio $z(\omega) = \log \frac{d\mu_1}{d\mu_0} \Big|_{\omega}$. We denote $Z_u = Z(W_u)$.

Definition 2.3. Private beliefs: *the private belief B_u of agent u is the posterior probability of the*

event $S = 1$ given u 's private signal W_u :

$$B_u = \mathbb{P}[S = 1|W_u].$$

Definition 2.4. Bounded and unbounded private beliefs: *In the framework of Definition 2.2 we say that the agents have **bounded private beliefs** if there exists an $\epsilon > 0$ such that $\mathbb{P}[\epsilon \leq B_u \leq 1 - \epsilon] = 1$.*

*Accordingly, private beliefs are **unbounded from below** when for every $\epsilon > 0$ it holds that $\mathbb{P}[B_u < \epsilon] > 0$, and **unbounded from above** when for every $\epsilon > 0$ it holds that $\mathbb{P}[B_u > 1 - \epsilon] > 0$.*

*Finally, **unbounded private beliefs** are such that are unbounded from both below and from above.*

Note that beliefs can be unbounded even though $\mathbb{P}[B_u = 0] = \mathbb{P}[B_u = 1] = 0$, i.e., beliefs are unbounded even when can come arbitrarily close to $\{0, 1\}$, although these values may never actually be realized.

Definition 2.5. Common knowledge of beliefs and actions: *Consider a set of agents with private signals, as defined in Definitions 2.1 and 2.2. The private signals may be bounded or unbounded.*

Assume that furthermore each agent u has some additional information, and denote by \mathcal{F}_u the σ -algebra of what is known to agent u . We assume that \mathcal{F}_u is a sub σ -algebra of $\sigma(W_1, \dots, W_n)$, so that each agent's additional information is measurable in the aggregation of the private signals.

*Let the **posterior belief** of agent u be given by $X_u = \mathbb{P}[S = 1|\mathcal{F}_u]$.*

*Let the **posterior optimal action set** of agent u , L_u , be the set of actions that is more likely, conditioned on \mathcal{F}_u , to be equal to the state of the world.*

$$L_u = \begin{cases} \{0\} & X_u < 1/2 \\ \{1\} & X_u > 1/2 \\ \{0, 1\} & X_u = 1/2. \end{cases}$$

Common knowledge of beliefs is said to occur when X_u is almost surely \mathcal{F}_w measurable for all $u, w \in V$.

Common knowledge of actions is said to occur when L_u is almost surely \mathcal{F}_w measurable for all $u, w \in V$.

2.2 Common beliefs model

In this section we discuss the case of common knowledge of beliefs. Using the definitions above, Aumann's Agreement Theorem implies the following:

Theorem (Aumann). *Let V be a set of agents with common knowledge of beliefs (Definition 2.5). Then there exists a random variable X such that almost surely $X_u = X$ for all $u \in V$.*

It follows that there exists a common optimal action set L such that almost surely $L_u = L$ for all u . Note that the converse of this theorem is trivially true: if all beliefs are almost surely equal then they are almost surely common knowledge.

Our main theorem for this section is the following.

Theorem 1. *Let V be a set of n agents with common knowledge of beliefs (Definition 2.5), so that there exists a common posterior belief X such that almost surely $X = X_u = \mathbb{P}[S = 1 | \mathcal{F}_u]$ for all $u \in V$. Then*

$$X = \mathbb{P}[S = 1 | W_1, \dots, W_n].$$

The extension of this theorem to private signals that are not identically distributed (but still conditionally independent) is straightforward; essentially the same proof applies.

2.2.1 Uncorrelated private signals

In this subsection we relax the conditional independence assumption on private signals, requiring instead that private signals are conditionally *uncorrelated*. Recall that $Z_u = \frac{d\mu_1}{d\mu_0}(W_u)$.

Definition 2.6. *Private signals are conditionally uncorrelated when $\text{Cov}[Z_u, Z_v | S] = 0$ for all $u, v \in V$.*

For example, when private signals are conditionally *pairwise* independent then they are uncorrelated.

Our main theorem here is the following.

Theorem 2. *Let V be a set of n agents with common knowledge of beliefs (Definition 2.5), so that there exists a common posterior belief X such that almost surely $X = X_u = \mathbb{P}[S = 1 | \mathcal{F}_u]$ for all $u \in V$. Let private signals be uncorrelated conditioned on S , but perhaps otherwise conditionally dependent.*

Then there exists a constant $D = D(\mu_0, \mu_1)$ depending only on μ_1 and μ_0 such that all agents take the optimal action with probability at least $1 - 4D/(n + D)$, where n is the number of agents:

$$\mathbb{P}[L = \{S\}] \geq 1 - \frac{4D}{n + D}.$$

If we impose no conditions at all on the dependence structure of the private signals, then it may be the case nothing at all may be learned, even when posteriors beliefs are common knowledge. Consider the following well known example (e.g., [8]).

Example 2.7. *Let the private signals be i.i.d. uniform bits (i.e., $\mathbb{P}[W_u = 1] = 1/2$), and let S equal the sum of the private signals, modulo 2.*

Clearly agent u 's private belief $\mathbb{P}[S = 1 | W_u]$ equals one half, regardless of W_u . Hence private beliefs are common knowledge, without need for any communication. However, the agents' action will only equal S w.p. half, independent of the number of agents.

The following next example shows that the bound of Theorem 2 is asymptotically tight.

Example 2.8. *Let V be the set of agents such that $|V| = n$ and n is divisible by 4. Let S be chosen uniformly at random from $\{0, 1\}$. Let U be a subset of V , chosen uniformly at random, and independent of S , from all the subsets of V of size $3n/4$. Let \hat{S} be an additional random variable such that*

$$\hat{S} = \begin{cases} S & \text{w.p. } p(n) \\ 1 - S & \text{w.p. } 1 - p(n), \end{cases}$$

where \hat{S} is independent of U and

$$p(n) = \frac{1}{2} + \frac{1}{2} \sqrt{1 - \frac{3}{n-1}}.$$

The private signals are:

$$W_u = \begin{cases} \hat{S} & \text{if } u \in U \\ 1 - \hat{S} & \text{otherwise} \end{cases}.$$

The communication protocol consists of all agents revealing their private signals to all other agents.

Now, it is easy to verify that private signals are pairwise-independent, conditioned on S :

$$\mathbb{P}[W_u = 1, W_w = 1 | S] = \mathbb{P}[W_u = 1 | S] \mathbb{P}[W_w = 1 | S],$$

and so in particular private signals are uncorrelated and this example satisfies the conditions of Theorem 2.

Note that

$$p(n) = 1 - \frac{3}{4} \cdot \frac{1}{n-1} + O\left(\frac{1}{n^2}\right),$$

so that asymptotically the probability that $\hat{S} \neq S$ is inversely proportional to n . Since the private signals are dependent on S only through \hat{S} , it follows that even though the agents gain complete knowledge of the private signals, they still may choose the wrong action with probability proportional to $1/n$.

Example 2.8 does not exclude the possibility that when signals are uncorrelated then information is still optimally aggregated, although this optimum may only reveal S with probability inversely proportional to the number of agents. This can be achieved by combining Examples 2.8 and 2.7.

Example 2.9. *Let each agents' private signal consist of two bits. The first bits of the agents are picked as in Example 2.7, while the second bits are pick as in Example 2.8.*

The communication protocol consists of all agents revealing their second bit to all other agents.

Here too the agents choose the action S with probability $1 - 1/n$. This is despite the fact that optimal aggregation of information would result in them choosing S with probability 1.

2.3 Common actions model

In this section we discuss the case of common knowledge of actions, as defined above in Definition 2.5.

As shown in Sebenius and Geanakoplos [19] (for finitely supported private signals) and more generally in Geanakoplos [9], common knowledge of optimal action sets implies that optimal actions are identical.

Theorem (Geanakoplos). *Consider a set of agents with common knowledge of actions (Definition 2.5). Then there exists an L such that almost surely $L_u = L$ for all $u \in V$.*

The boundedness of beliefs (Definition 2.4) play an important role in the case of common actions. When private beliefs are bounded then common knowledge of optimal action sets does not imply asymptotic learning, as shown by the following example³. However, when private beliefs are unbounded then asymptotic learning occurs, as we show below.

Example 2.10. *Let there be $n > 100$ agents, and call the first hundred “the Senate”. The private signal is a bit that is independently equal to S with probability $2/3$. At stage one, each senator gains knowledge of the private signals of all the other senators. At stage two, all agents gain knowledge of this optimal action of the Senate.*

It is easy to convince oneself that after stage one, all senators adopt the same action, which depends only on their senatorial signals. Also, it is easy to see that after at stage two, by Bayes’ Law, the rest of the agents disregard their own private signals and adopt the Senate’s optimal action set. Hence all the agents’ optimal action sets are common knowledge after this stage.

The probability that the Senate makes a mistake is strictly constant and does not depend on the number of agents n . Hence the probability that the agents choose the wrong action does *not* tend to zero as n tends to infinity, even though this action set is common knowledge.

We show that this cannot be the case when private beliefs are unbounded, in the following theorem which is the main result of this section.

Theorem 3. *Let V be a set of n agents with common knowledge of actions (Definition 2.5), so that there exists a common posterior actions set L such that almost surely $L = L_u$ for all $u \in V$. Assume that beliefs unbounded from below (Definition 2.4).*

Then there exists a sequence $q(n) = q(n, \mu_0, \mu_1)$, depending only μ_1 and μ_0 , such that $q(n) \rightarrow 1$ as $n \rightarrow \infty$, and when $S = 0$ then the agents take action “0” with probability at least $q(n)$, where n is the number of agents:

$$\mathbb{P}[L = \{0\} | S = 0] \geq q(n).$$

By the symmetry of the states of the world, an immediate consequence of this theorem is that asymptotic learning occurs when L is common knowledge and private signals are unbounded in both directions:

$$\mathbb{P}[L = \{S\}] \geq q(n)$$

for some $q(n) \rightarrow 1$, when private beliefs are unbounded.

3 Proofs

3.1 Common beliefs model

3.1.1 Preliminaries

We recall the details of the common beliefs model. S is binary state of the world picked uniformly at random. The information available to each agent u , denoted by \mathcal{F}_u , includes a private signal

³This example is reminiscent of Bala and Goyal’s [2] *royal family* graph.

W_u , picked from either μ_1 or μ_0 , depending on whether $S = +$ or $S = -$. Agent u 's posterior belief is $X_u = \mathbb{P}[S = 1|\mathcal{F}_u]$, and its posterior action is L_u is given by

$$L_u = \begin{cases} \{0\} & X_u < 1/2 \\ \{1\} & X_u > 1/2 \\ \{0, 1\} & X_u = 1/2. \end{cases}$$

Beliefs are common knowledge: X_w is \mathcal{F}_u measurable for all $w \in V$. By Aumann's Agreement Theorem [1], there exists an X such that $X_u = X$ for all u . Hence there also exists an L such that $L_u = L$ for all u .

Our main result for this model is the following.

Theorem. *1 Let V be a set of n agents with common knowledge of beliefs (Definition 2.5), so that there exists a common posterior belief X such that almost surely $X = X_u = \mathbb{P}[S = 1|\mathcal{F}_u]$ for all $u \in V$. Then*

$$X = \mathbb{P}[S = 1|W_1, \dots, W_n].$$

Proof. Recall that Z_u , agent u 's private log-likelihood ratio, is

$$Z_u = \log \frac{d\mu_1}{d\mu_0} \Big|_{W_u} = \log \frac{\mathbb{P}[S = 1|W_u]}{\mathbb{P}[S = 0|W_u]}.$$

Denote $Z = \sum_u Z_u$. Then since the private signals are conditionally independent it follows by Bayes' rule that

$$\mathbb{P}[S = 1|W_1, \dots, W_n] = \text{logit}(Z), \tag{1}$$

where $\text{logit}(z) = e^z/(e^z + e^{-z})$. Since

$$X = \mathbb{P}[S = 1|X] = \mathbb{E}[\mathbb{P}[S = 1|X, W_1, \dots, W_n]|X]$$

then

$$X = \mathbb{E}[\text{logit}(Z)|X], \tag{2}$$

since, given the private signals (W_1, \dots, W_n) , further conditioning on X (which is a function of the private signals) does not change the probability of the event $S = 1$.

Our goal is to show that $X = \mathbb{P}[S = 1|W_1, \dots, W_n]$. We will do this by showing that conditioned on X , Z and $\text{logit}(Z)$ are linearly independent. It will follow that conditioned on X , Z is constant, so that $Z = Z(X)$ and

$$X = \mathbb{P}[S = 1|X] = \mathbb{P}[S = 1|Z(X)] = \mathbb{P}[S = 1|W_1, \dots, W_n].$$

By the law of total expectation we have that

$$\mathbb{E}[Z_u \cdot \text{logit}(Z)|X] = \mathbb{E}[\mathbb{E}[Z_u \text{logit}(Z)|X, Z_u]|X].$$

Note that $\mathbb{E}[Z_u \text{logit}(Z)|X, Z_u] = Z_u \mathbb{E}[\text{logit}(Z)|X, Z_u]$ and so we can write

$$\mathbb{E}[Z_u \cdot \text{logit}(Z)|X] = \mathbb{E}[Z_u \mathbb{E}[\text{logit}(Z)|X, Z_u]|X].$$

Since Z_u is \mathcal{F}_u measurable, and since, by Eq. 2, $X = \mathbb{E}[\text{logit}(Z)|\mathcal{F}_u] = \mathbb{E}[\text{logit}(Z)|X]$, then $X = \mathbb{E}[\text{logit}(Z)|X, Z_u]$ and so it follows that

$$\mathbb{E}[Z_u \cdot \text{logit}(Z)|X] = \mathbb{E}[Z_u X|X] = X \mathbb{E}[Z_u|X] = \mathbb{E}[\text{logit}(Z)|X] \mathbb{E}[Z_u X|X]. \quad (3)$$

where the last equality is another substitution of Eq. 2. Summing this equation (3) over $u \in V$ we get that

$$\mathbb{E}[Z \cdot \text{logit}(Z)|X] = \mathbb{E}[\text{logit}(Z)|X] \mathbb{E}[Z|X]. \quad (4)$$

Now, since $\text{logit}(Z)$ is a monotone function of Z , by Chebyshev's sum inequality we have that

$$\mathbb{E}[Z \cdot \text{logit}(Z)|X] \geq \mathbb{E}[\text{logit}(Z)|X] \mathbb{E}[Z|X] \quad (5)$$

with equality only if Z (or, equivalently $\text{logit}(Z)$) is constant. Hence Z is constant conditioned on X and the proof is concluded. \square

3.1.2 Uncorrelated private signals

We now turn to prove Theorem 2. We here relax the conditional independence condition on private signals and replace it with the weaker condition of uncorrelated private signals (Definition 2.6).

When private signals are uncorrelated then Theorem 2 states that all the agents pick the correct action $L = \{S\}$ with high probability: $\mathbb{P}[L = \{S\}] \geq 1 - \frac{4D}{n+D}$. The constant D here depends on how informative the private signals are. We call it the **noise to signal ratio** of (μ_1, μ_0) , and define it as follows:

Definition 3.1. Let μ_1, μ_0 be two distributions on (Ω, \mathcal{O}) , such that the Radon-Nikodym derivative $z(\omega) = \frac{d\mu_1}{d\mu_0} \Big|_{\omega}$ exists everywhere. Then the **noise to signal ratio** of the pair (μ_1, μ_0) is

$$D = 2 \frac{\text{Var}_{\mu_1} Z + \text{Var}_{\mu_0} Z}{(D_{KL}(\mu_1||\mu_0) + D_{KL}(\mu_1||\mu_0))^2}. \quad (6)$$

Where $D_{KL}(\cdot||\cdot)$ denotes the Kullback-Leibler divergence [11].

The reader may object that any of the four quantities involved in the definition of D may diverge. We proceed with this section assuming this is not the case, and address this technical point below in subsection 3.1.3.

Recall the definition $X_u = \mathbb{P}[S = 1|\mathcal{F}_u]$. Since $S \in \{0, 1\}$, this is equivalent to $X_u = \mathbb{E}[S|\mathcal{F}_u]$, which in turn makes X_u the minimum variance unbiased estimator (or MVUE, c.f. [4]) of S given \mathcal{F}_u :

$$X_u = \underset{M \in \mathcal{F}_u, \mathbb{E}[M] = \mathbb{E}[S]}{\text{argmin}} \text{Var}[M - S],$$

Recalling that there exists a random variable X such that almost surely $X_u = X$ we can likewise write $X = \underset{M \in \mathcal{F}_u}{\text{argmin}} \text{Var}[M - S]$, for all $u \in V$.

We show that for large networks X is indeed a good estimator of S , in the following sense.

Theorem 3.2. Consider the common beliefs model (Definition 2.5) with agent set V such that $|V| = n$ and with uncorrelated private signals (Definition 2.6). Then

$$\text{Var}[X - S] \leq \frac{D}{n + D}.$$

where the constant $D = D(\mu_1, \mu_0)$ is the noise to signal ratio of (μ_1, μ_0) .

Our main result for this model is an easy consequence of this theorem:

Theorem (2). Let V be a set of n agents with common knowledge of beliefs (Definition 2.5), so that there exists a common posterior belief X such that almost surely $X = X_u = \mathbb{P}[S = 1 | \mathcal{F}_u]$ for all $u \in V$. Let private signals be uncorrelated conditioned on S , but perhaps otherwise conditionally dependent.

Then there exists a constant $D = D(\mu_0, \mu_1)$ depending only on μ_1 and μ_0 such that all agents take the optimal action with probability at least $1 - 4D/(n + D)$, where n is the number of agents:

$$\mathbb{P}[L = \{S\}] \geq 1 - \frac{4D}{n + D}.$$

Proof. Since X is the minimum variance unbiased estimator of S given \mathcal{F}_u , it is in particular unbiased: $\mathbb{E}[X] = \mathbb{E}[S]$. Hence we can strengthen the statement of Theorem 3.2, and in fact write

$$\mathbb{E}[(X - S)^2] \leq \frac{D}{n + D}. \quad (7)$$

By Markov's inequality we have then that

$$\mathbb{P}[|X - S| \geq \frac{1}{2}] \leq \frac{4D}{n + D}.$$

Since $L_u = \{S\}$ for all u whenever $|X - S| < \frac{1}{2}$ then the theorem follows. \square

To prove Theorem 3.2 we shall need a few standard lemmas:

Lemma 3.3. If $\mu_1 \neq \mu_0$ then $\text{Cov}[S, Z_u] > 0$.

Proof. Let $D_{KL}(\cdot || \cdot)$ denote the Kullback-Leibler divergence [11]. Then

$$\begin{aligned} \text{Cov}[S, Z_u] &= \mathbb{E}[SZ_u] - \mathbb{E}[S]\mathbb{E}[Z_u] \\ &= \frac{1}{4}\mathbb{E}[Z_u|S = 1] - \frac{1}{4}\mathbb{E}[Z_u|S = 0] \\ &= \frac{1}{4} \int \log \frac{d\mu_1}{d\mu_0} (d\mu_1 - d\mu_0) \\ &= \frac{1}{4} [D_{KL}(\mu_1 || \mu_0) + D_{KL}(\mu_0 || \mu_1)] \end{aligned}$$

The lemma follows from the fact that KL-divergence is zero only for identical distributions and positive otherwise. \square

Since the private beliefs are identically distributed then $\text{Cov}[S, Z_u]$ is independent of u .

An equivalent statement to the lemma above is that $\mathbb{E}[Z_u|S=1] > \mathbb{E}[Z_u|S=0]$. Since the private signals are uncorrelated conditioned on S , we expect that for large n their average will be close to either $\mathbb{E}[Z_u|S=1]$ if $S=1$ or $\mathbb{E}[Z_u|S=0]$ if $S=0$. Hence we can expect that

$$Y = \frac{1}{n} \sum_u \frac{Z_u - \mathbb{E}[Z_u|S=0]}{\mathbb{E}[Z_u|S=1] - \mathbb{E}[Z_u|S=0]} \quad (8)$$

will be a good estimator for S . To prove our theorem we show in lemma 3.4 that this is indeed the case. We will finish the proof by furthermore showing that the agents' limiting estimator X is as good as Y .

Lemma 3.4. $\text{Var}[Y - S] = \frac{D}{4n}$, where $D = D(\mu_1, \mu_0)$.

Proof. From the definition of Y it follows that $\mathbb{E}[Y] = \mathbb{E}[S]$. Hence

$$\text{Var}[Y - S] = \mathbb{P}[S=1] \mathbb{E}[(Y - S)^2|S=1] + \mathbb{P}[S=0] \mathbb{E}[(Y - S)^2|S=0].$$

By the definition of Y this means that

$$\text{Var}[Y - S] = \frac{1}{n^2} \cdot \frac{\text{Var}[\sum_u Z_u|S=1] + \text{Var}[\sum_u Z_u|S=0]}{2(\mathbb{E}[Z_u|S=1] - \mathbb{E}[Z_u|S=0])^2},$$

and since the private belief log-likelihood ratios Z_u are uncorrelated conditioned on S , we have that

$$\text{Var}\left[\sum_u Z_u \middle| S=s\right] = \sum_u \text{Var}[Z_u|S=s] = n \text{Var}[Z_u|S=s]$$

so that

$$\text{Var}[Y - S] = \frac{D}{4n},$$

where $D = D(\mu_1, \mu_0)$ is the noise to signal ratio of Definition 3.1 above. \square

We would next like to bound the variance of Y . By the triangle inequality it follows from lemma 3.4 above and from the fact that $\text{Var}[S] = 1/4$ that $\text{Var}[Y] \leq 1/4 + O(\sqrt{D/n})$. We show that in fact $\text{Var}[Y] = 1/4 \cdot (1 + D/n)$.

Lemma 3.5. $\text{Cov}[S, Y] = 1/4$ and $\text{Var}[Y] = \frac{1}{4} \cdot (1 + \frac{D}{n})$.

Proof. We first note that

$$\begin{aligned} \text{Cov}[Z_u, S] &= \mathbb{E}[Z_u S] - \mathbb{E}[S] \mathbb{E}[Z_u] \\ &= \frac{1}{2} \mathbb{E}[Z_u|S=1] - \frac{1}{2} (\frac{1}{2} \mathbb{E}[Z_u|S=1] + \frac{1}{2} \mathbb{E}[Z_u|S=0]) \\ &= \frac{1}{4} (\mathbb{E}[Z_u|S=1] - \mathbb{E}[Z_u|S=0]) \end{aligned}$$

Since the different Z_u 's are independent conditioned on S it follow from the definition of Y (Eq. 8) that $\text{Cov}[S, Y] = 1/4$. The claim $\text{Var}[Y] = \frac{1}{4} \cdot (1 + \frac{D}{n})$ follows by substituting $\text{Var}[Y - S] = D/(4n)$, the result of lemma 3.4 above, into the identity $\text{Var}[Y - S] = \text{Var}[S] + \text{Var}[Y] - 2\text{Cov}[S, Y]$. \square

We are now ready to prove Theorem 3.2:

Proof. Recall that Z_u is u 's private belief and as such is in \mathcal{F}_u , as is $X + bZ_u$, for any $b \in \mathbb{R}$. Therefore, since $X = X_u$ is the minimum variance unbiased estimator of S given \mathcal{F}_u , for all $b \in \mathbb{R}$

$$\text{Var}[X - S] \leq \text{Var}[(X + bZ_u) - S].$$

We therefore have that

$$\left. \frac{\partial \text{Var}[(X + bZ_u) - S]}{\partial b} \right|_{b=0} = 0, \quad (9)$$

or that $\text{Cov}[X, Z_u] = \text{Cov}[S, Z_u]$. Since this is true for all u , by the bilinearity of covariance we have that

$$\text{Cov}[X, Y] = \text{Cov}[S, Y]. \quad (10)$$

Consider the set of zero expectation random variables measurable in our measure space which have a second moment. This set is a real Hilbert space with $\text{Cov}[\cdot, \cdot]$ as its inner product. Hence we can decompose $(S - \mathbb{E}[S])$ (a zero expectation r.v.) to the following sum:

$$S - \mathbb{E}[S] = \alpha(Y - \mathbb{E}[Y]) + W_S, \quad (11)$$

with $\alpha \in \mathbb{R}$ and $0 = \text{Cov}[Y - \mathbb{E}[Y], W_S] = \text{Cov}[Y, W_S]$. It follows that

$$\text{Cov}[S, Y] = \text{Cov}[\alpha Y + W_S, Y] = \alpha \text{Var}[Y]. \quad (12)$$

We can likewise decompose

$$X - \mathbb{E}[X] = \beta(Y - \mathbb{E}[Y]) + W_X,$$

with $\beta \in \mathbb{R}$ and $\text{Cov}[Y, W_X] = 0$. But since $\text{Cov}[X, Y] = \text{Cov}[S, Y]$ (Eq. 10) it must be that in fact $\beta = \alpha$ and we can write

$$X - \mathbb{E}[X] = \alpha(Y - \mathbb{E}[Y]) + W_X. \quad (13)$$

Since $\mathbb{E}[X] = \mathbb{E}[S]$ we can subtract Eq. 11 from Eq. 13 and get $X - S = W_X - W_S$. We will show that $\text{Var}[X - S]$ is small by showing the both $\text{Var}[W_S]$ and $\text{Var}[W_X]$ are small.

Since $\text{Cov}[Y, W_S] = 0$ and also $\text{Cov}[Y, W_X] = 0$ we can take the variance of Eq. 11 and Eq. 13 to arrive at:

$$\begin{aligned} \text{Var}[S] &= \alpha^2 \text{Var}[Y] + \text{Var}[W_S] \\ \text{Var}[X] &= \alpha^2 \text{Var}[Y] + \text{Var}[W_X]. \end{aligned}$$

Now, by Eq. (12) we know that $\alpha = \text{Cov}[S, Y] / \text{Var}[Y]$, and lemma 3.5 states that $\text{Cov}[S, Y] = 1/4$ and $\text{Var}[Y] = \frac{1}{4} \left(1 + \frac{D}{4n}\right)$. Hence

$$\alpha^2 \text{Var}[Y] = \frac{\text{Cov}[S, Y]^2}{\text{Var}[Y]} = \frac{1}{4} \cdot \frac{1}{1 + D/n}$$

and so

$$\begin{aligned}\text{Var}[S] &= \frac{1}{4} \cdot \frac{1}{1 + D/n} + \text{Var}[W_S] \\ \text{Var}[X] &= \frac{1}{4} \cdot \frac{1}{1 + D/n} + \text{Var}[W_X].\end{aligned}$$

Now $\text{Var}[S] = 1/4$ and $\text{Var}[X] \leq 1/4$, since X is bounded between 0 and 1. Hence

$$\begin{aligned}\text{Var}[W_S] &= \frac{1}{4} \cdot \frac{D}{n + D} \\ \text{Var}[W_X] &\leq \frac{1}{4} \cdot \frac{D}{n + D}.\end{aligned}$$

We have then that

$$\begin{aligned}\text{Var}[X - S] &= \text{Var}[W_S - W_X] \\ &\leq 2(\text{Var}[W_S] + \text{Var}[W_X]) \\ &\leq \frac{D}{n + D}\end{aligned}$$

□

3.1.3 Note on private signals with infinite second moments

The proof of Theorem 3.2 above applies when Z has a second moment. We comment here on how it may be extended to the case that the second moment of Z_u diverges.

In this case the difficulty that arises is that D and Y are not well defined. To mend this we substitute Y_u^M for Z_u in the definitions of Y (Eq. 8) and D (Eq. 6) and in the proofs that follow. We fix M and take Y_u^M to equal Z_u when $|Z_u| < M$ and $Y_u^M = 0$ otherwise, choosing M large enough so that Y_u^M is an informative signal with respect to S ; let $\mu_{1,M}$ and $\mu_{0,M}$ be the distributions of Y_u^M when $S = 1$ and $S = 0$, respectively. Then we choose M that enough so that $\mu_{1,M} \neq \mu_{0,M}$.

Now, Y_u^M is bounded and therefore has both a first and a second moment, which our proof requires. Furthermore, Y_u is measurable by agent u ; substituting it for Z_u is tantamount to having agent u ignore its private signal when it is too strong. The proofs of the lemmas and the theorem remain valid, with a small variation to the definition of the noise to signal ratio D - Theorems 1 and 3.2 hold with $D = D(\mu_{1,M}, \mu_{0,M})$, rather than $D(\mu_1, \mu_0)$, which isn't well defined in this case.

3.2 Common actions model

We recall the details of the common actions model (Definition 2.5). S is binary state of the world picked uniformly at random. The information available to each agent u , denoted by \mathcal{F}_u , includes a

private signal W_u , picked from either μ_1 or μ_0 , depending on whether $S = +$ or $S = -$. Agent u 's posterior belief is $X_u = \mathbb{P}[S = 1|\mathcal{F}_u]$, and its posterior optimal action set is L_u is given by

$$L_u = \begin{cases} \{0\} & X_u < 1/2 \\ \{1\} & X_u > 1/2 \\ \{0, 1\} & X_u = 1/2. \end{cases}$$

By virtue of it being a common actions model, \mathcal{F}_u further includes L_w for all $w \in V$. By Geanakoplos [9], there exists an L such that $L_u = L$ for all u .

Definition 3.6. We denote by B_u the probability of $S = 1$ given agent u 's private signal:

$$B_u = \mathbb{P}[S = 1|W_u].$$

We consider here the case of unbounded beliefs (Definition 2.4), so that for any $\epsilon > 0$ it holds that $\mathbb{P}[B_u < \epsilon] > 0$ and $\mathbb{P}[B_u > 1 - \epsilon] > 0$.

Our main result for this model is the following:

Theorem 3.7 (3). Let V be a set of n agents with common knowledge of actions (Definition 2.5), so that there exists a common posterior actions set L such that almost surely $L = L_u$ for all $u \in V$. Assume that beliefs unbounded from below (Definition 2.4).

Then there exists a sequence $q(n) = q(n, \mu_0, \mu_1)$, depending only μ_1 and μ_0 , such that $q(n) \rightarrow 1$ as $n \rightarrow \infty$, and when $S = 0$ then the agents take action "0" with probability at least $q(n)$, where n is the number of agents:

$$\mathbb{P}[L = \{0\}|S = 0] \geq q(n).$$

This theorem follows trivially from the following theorem which we prove below:

Theorem 3.8. In the common actions model, fix μ_1 and μ_0 . Consider a sequence of agent sets $\{V_n\}_{n=1}^\infty$, with $|V_n| = n$. Denote the common optimal action set of the agents in V_n by L_n . Then if private signals are unbounded from below then

$$\lim_{n \rightarrow \infty} \mathbb{P}[L_n = \{0\}|S = 0] = 1.$$

We shall need two standard lemmas to prove this theorem.

Lemma 3.9. $\mathbb{P}[S = 1|B_u] = B_u$.

Proof. Since B_u is a function of W_u then

$$\begin{aligned} \mathbb{P}[S = 1|B_u = b_u] &= \mathbb{E}[\mathbb{P}[S = 1|W_u]|B_u(W_u) = b_u] \\ &= \mathbb{E}[B_u|B_u = b_u] \\ &= b_u. \end{aligned}$$

□

The two following corollary follows trivially from this lemma:

Corollary 3.10. $\mathbb{P}[S = 0|B_u < \epsilon] > 1 - \epsilon$.

Lemma 3.11 is a version of Chebyshev's inequality, quantifying the idea that the expectation of a random variable Z , conditioned on some event A , cannot be much lower than its unconditional expectation when A has high probability.

Lemma 3.11. *Let Z be a real valued random variable with finite variance, and let A be an event. Then*

$$\mathbb{E}[Z] - \sqrt{\frac{\text{Var}[Z]}{\mathbb{P}[A]}} \leq \mathbb{E}[Z|A] \leq \mathbb{E}[Z] + \sqrt{\frac{\text{Var}[Z]}{\mathbb{P}[A]}}$$

Proof. By Cauchy-Schwarz

$$|\mathbb{E}[Z\mathbf{1}(A)] - \mathbb{E}[Z]\mathbb{P}[A]| = |\mathbb{E}[(Z - \mathbb{E}[Z]) \cdot \mathbf{1}(A)]| \leq \sqrt{\text{Var}[Z]\mathbb{P}[A]}.$$

Dividing by $\mathbb{P}[A]$ and noting that $\mathbb{E}[Z\mathbf{1}(A)]/\mathbb{P}[A] = \mathbb{E}[Z|A]$ we obtain the statement of the lemma. \square

We are now ready to prove Theorem 3.8.

Proof of Theorem 3.8. For simplicity of notation we write " $L_n = 0$ " rather than " $L_n = \{0\}$ ", in this proof.

We would like to show that $\lim_{n \rightarrow \infty} \mathbb{P}[L_n = 0|S = 0] = 1$. For ease of notation we will denote $q_n = \mathbb{P}[L_n \neq 0|S = 0]$ and show that $\lim_{n \rightarrow \infty} q_n = 0$.

Pick an arbitrary agent u . Recall that by definition $X_u = \mathbb{P}[S = 1|\mathcal{F}_u]$. By lemma 3.2 above $\mathbb{P}[S = 1|B_u] = B_u$, and so

$$\mathbb{E}[X_u|B_u] = \mathbb{E}[\mathbb{P}[S = 1|\mathcal{F}_u]|B_u] = \mathbb{P}[S = 1|B_u] = B_u.$$

Applying Markov's inequality to X_u we have that $\mathbb{P}[X_u \geq \frac{1}{2}|B_u < \epsilon] < 2\epsilon$, and in particular $\mathbb{P}[X_u \geq \frac{1}{2}, S = 0|B_u < \epsilon] < 2\epsilon$.

Since all agents have the same optimal action set L_n , the choice of u is immaterial and the event " $X_u \geq \frac{1}{2}$ " is, by definition of L_n , the same event as " $L_n \neq 0$ ". Hence

$$\mathbb{P}[L_n \neq 0, S = 0|B_u < \epsilon] < 2\epsilon.$$

By application of Bayes' Law we have that

$$\mathbb{P}[B_u < \epsilon|L_n \neq 0, S = 0] < 2\epsilon \frac{\mathbb{P}[B_u < \epsilon]}{\mathbb{P}[L_n \neq 0, S = 0]}.$$

Recalling our definition $q_n = \mathbb{P}[L_n \neq 0|S = 0]$, we have that $\mathbb{P}[L_n \neq 0, S = 0] = \mathbb{P}[S = 0]q_n = \frac{1}{2}q_n$ so that we can write

$$\mathbb{P}[B_u < \epsilon|L_n \neq 0, S = 0] < 4\epsilon \frac{\mathbb{P}[B_u < \epsilon]}{q_n}. \tag{14}$$

If we denote

$$K_n = \frac{1}{n} \sum_{u \in V} \mathbf{1}(B_u < \epsilon) \tag{15}$$

then by averaging Eq. 14 over all $u \in V$ we get that

$$\mathbb{E}[K_n | L_n \neq 0, S = 0] < 4\epsilon \frac{\mathbb{P}[B_u < \epsilon]}{q_n}. \quad (16)$$

We thus bound $\mathbb{E}[K_n | L_n \neq 0, S = 0]$ from above. We will show that q_n is small by also bounding it from above.

Applying lemma 3.11 to K_n and the event “ $L_n \neq 0$ ” (under the conditional measure $S = 0$), this lemma yields that

$$\mathbb{E}[K_n | L_n \neq 0, S = 0] \geq \mathbb{E}[K_n | S = 0] - \sqrt{\frac{\text{Var}[K_n | S = 0]}{q_n}}.$$

Since the agents’ private signals (and hence their private beliefs) are independent conditioned on $S = 0$, K_n (conditioned on S) is the average of n i.i.d. variables. Hence $\text{Var}[K_n | S = 0] = n^{-1} \text{Var}[\mathbf{1}(B_u < \epsilon) | S = 0]$ and $\mathbb{E}[K_n | S = 0] = \mathbb{P}[B_u < \epsilon | S = 0]$. Thus we have that

$$\mathbb{E}[K_n | L_n \neq 0, S = 0] \geq \mathbb{P}[B_u < \epsilon | S = 0] - n^{-1/2} \sqrt{\frac{\text{Var}[\mathbf{1}(B_u < \epsilon) | S = 0]}{q_n}}. \quad (17)$$

Joining this bound (Eq. 17) with the bound from Eq. 16, we get that

$$\mathbb{P}[B_u < \epsilon | S = 0] - n^{-1/2} \sqrt{\frac{\text{Var}[\mathbf{1}(B_u < \epsilon) | S = 0]}{q_n}} < 2\epsilon \frac{\mathbb{P}[B_u < \epsilon]}{q_n}$$

or

$$q_n < \frac{2\epsilon \mathbb{P}[B_u < \epsilon] + n^{-1/2} \sqrt{q_n \text{Var}[\mathbf{1}(B_u < \epsilon) | S = 0]}}{\mathbb{P}[B_u < \epsilon | S = 0]}. \quad (18)$$

Here we used the theorem hypothesis that $\mathbb{P}[B_u < \epsilon | S = 0] > 0$, or that the signals are unbounded from below.

Taking the limit $n \rightarrow \infty$ we get that

$$\lim_{n \rightarrow \infty} q_n < 2\epsilon \frac{\mathbb{P}[B_u < \epsilon]}{\mathbb{P}[B_u < \epsilon | S = 0]},$$

which by application of Bayes’ law means that

$$\lim_{n \rightarrow \infty} q_n < \frac{\epsilon}{\mathbb{P}[S = 0 | B_u < \epsilon]},$$

and since by corollary 3.10 above we know that $\mathbb{P}[S = 0 | B_u < \epsilon] > 1 - \epsilon$, then

$$\lim_{n \rightarrow \infty} q_n < \frac{\epsilon}{1 - \epsilon}.$$

Since this holds for all ϵ , we have shown that $\lim_{n \rightarrow \infty} q_n = 0$. □

The proof above in fact implies a more quantitative bound for q_n in terms of n and the private belief distribution:

Corollary 3.12.

$$q_n \leq \min_{\epsilon > 0} \max \left\{ \frac{2\epsilon}{1-\epsilon}, \frac{4}{n\mathbb{P}[B_u < \epsilon | S = 0]} \right\}. \quad (19)$$

Proof. In Eq. 18 above, we consider two case, depending on which of the first or the second term in the denominator of the r.h.s. is larger, and using the fact that $a + b \leq 2 \max\{a, b\}$.

In the case that $2\epsilon\mathbb{P}[B_u < \epsilon]$ is the larger term then the same derivation of the proof applies so that

$$q_n \leq 2 \frac{\epsilon}{1-\epsilon}.$$

In the second case we obtain

$$q_n \leq \frac{2n^{-1/2} \sqrt{q_n \text{Var}[\mathbf{1}(B_u < \epsilon) | S = 0]}}{\mathbb{P}[B_u < \epsilon | S = 0]},$$

which by squaring and simplifying yields

$$q_n \leq \frac{4}{n} \cdot \frac{\text{Var}[\mathbf{1}(B_u < \epsilon) | S = 0]}{\mathbb{P}[B_u < \epsilon | S = 0]^2} = \frac{4}{n} \cdot \frac{1 - \mathbb{P}[B_u < \epsilon | S = 0]}{\mathbb{P}[B_u < \epsilon | S = 0]} \leq \frac{4}{n} \cdot \frac{1}{\mathbb{P}[B_u < \epsilon | S = 0]}.$$

□

References

- [1] R. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- [2] V. Bala and S. Goyal. Learning from neighbours. *Review of Economic Studies*, 65(3):595–621, July 1998.
- [3] A. V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- [4] P. Bickel and K. Doksum. *Mathematical statistics: Basic ideas and selected topics*. 2d. ed. vol. 1 prentice hall. *Upper Saddle River, NJ*, 2001.
- [5] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascade. *Journal of Political Economy*, 100(5):992–1026, 1992.
- [6] P. DeMarzo, D. Vayanos, and J. Zwiebel. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118:909–968, 2003.
- [7] G. Ellison and D. Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 110(1):93–126, 1995.
- [8] J. Feigenbaum, L. Fortnow, D. Pennock, and R. Sami. Computation in a distributed information market. *Theoretical Computer Science*, 343(1-2):114–132, 2005.

- [9] J. Geanakoplos. Common knowledge. In *Proceedings of the 4th conference on Theoretical aspects of reasoning about knowledge*, pages 254–315. Morgan Kaufmann Publishers Inc., 1992.
- [10] J. Geanakoplos and H. Polemarchakis. We can’t disagree forever* 1. *Journal of Economic Theory*, 28(1):192–200, 1982.
- [11] S. Kullback and R. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [12] R. McKelvey and T. Page. Common knowledge, consensus, and aggregate information. *Econometrica: Journal of the Econometric Society*, pages 109–127, 1986.
- [13] L. Ménager. Consensus, communication and knowledge: an extension with bayesian agents. *Mathematical Social Sciences*, 51(3):274–279, 2006.
- [14] D. Minehart and S. Scotchmer. Ex post regret and the decentralized sharing of information. *Games and Economic Behavior*, 27(1):114–131, 1999.
- [15] E. Mossel and O. Tamuz. Efficient bayesian learning in social networks with gaussian estimators. Preprint at <http://arxiv.org/abs/1002.0747>, 2010.
- [16] M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 253–254. ACM, 2009.
- [17] R. Parikh and P. Krasucki. Communication, consensus, and knowledge* 1. *Journal of Economic Theory*, 52(1):178–189, 1990.
- [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] L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.