

Aggregation and Influence in Teams of Imperfect Decision Makers

by

Joong Bum Rhim

B.S. Electrical Engineering
KAIST, 2008

S.M., Electrical Engineering and Computer Science
Massachusetts Institute of Technology, 2010

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

© Massachusetts Institute of Technology 2014. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 21, 2014

Certified by.....
Vivek K Goyal
Visiting Scientist, Research Laboratory of Electronics
Assistant Professor of Electrical and Computer Engineering, Boston University
Thesis Supervisor

Accepted by.....
Professor Leslie A. Kolodziejcki
Chair, Committee on Graduate Students

Aggregation and Influence in Teams of Imperfect Decision Makers

by

Joong Bum Rhim

Submitted to the Department of Electrical Engineering and Computer Science
on May 21, 2014, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

Bayesian hypothesis testing inevitably requires prior probabilities of hypotheses. Motivated by human decision makers, this thesis studies how binary decision making is performed when the decision-making agents use imperfect prior probabilities. Three detection models with multiple agents are investigated: distributed detection with symmetric fusion, sequential detection with social learning, and distributed detection with symmetric fusion and social learning.

In the distributed detection with symmetric fusion, we consider the agents to be a team aiming to minimize the Bayes risk of the team's decision. In this model, incorrect beliefs reduce the chance of the agents from being right so always lead to an increase in the Bayes risk of the decision-making team.

In contrast, the role of beliefs is more complicated in the sequential detection model with social learning, where agents observe public signals, which are decisions made by other agents. Since each agent affects the minimum possible Bayes risk for subsequent agents, she may have a mixed objective including her own Bayes risk and the Bayes risks of subsequent agents. For an earlier-acting agent, it is shown that being informative to later-acting agents is different from being right. When private signals are described by Gaussian likelihoods, informative earlier-acting agents should be open-minded toward the unlikely hypothesis. Social learning helps imperfect agents who have favorable incorrect beliefs outperform perfect agents who have correct beliefs.

Compared to in the sequential detection model, social learning is less influential in the distributed detection model with symmetric fusion. This is because social learning induces the evolution of the fusion rule in the distributed detection model, which countervails the other effect of social learning—belief update. In particular, social learning is futile when the agents observe conditionally independent and identically distributed private signals or when the agents require unanimity to make a decision. Since social learning is ineffective, imperfect agents cannot outperform perfect agents, unlike in the sequential detection model.

Experiments about human behavior were performed in team decision-making sit-

uations when people should optimally ignore public signals. The experiments suggest that when people vote with equal qualities of information, the ballots should be secret.

Thesis Supervisor: Vivek K Goyal

Title: Visiting Scientist, Research Laboratory of Electronics

Assistant Professor of Electrical and Computer Engineering, Boston University

Acknowledgments

As I recall my last six years at MIT, what I have learned mostly came from Vivek Goyal and the Signal Transportation and Information Representation (STIR) group. Vivek was the best teacher and advisor whom I can ever meet; he showed me a way to research problems, waited for me to think by myself, and taught me through lots of free discussion. The members of the STIR group, Lav Varshney, Daniel Weller, John Sun, Ahmed Kirmani, Andrea Colaço, Vahid Montazerhodjat, Donggeek Shin, and an honorary member Da Wang, also showed many virtues of a good researcher, which include showing earnestness, having intuition and creativity, and most importantly, enjoying their lives. Thanks to them, I am full of great memories at MIT. It is my honor and fortune to be a part of the STIR group.

I owe many thanks to my dissertation committee, Devavrat Shah, Lav Varshney, and Juanjuan Zhang. As my thesis is not about orthodox topics in signal processing, advice from their various experiences and interests helped me a lot expanding my horizons and developing this thesis. I would like to thank my other mentor, Alan Oppenheim. He was always concerned for me and gave counsel to help me being comfortable and happy at MIT.

I want to mention the support from the Korea Foundation for Advanced Studies (KFAS). Not only did they financially support my study, but also cared about my student life.

In addition to the research, I also enjoyed collaboration with Jesse Weinstein-Gould and Wynn Sullivan of Oklahoma City Thunder. It provided me with a bridge from theory to practice and inspired me in choosing my career after graduation.

I would like to thank my friends, too. They are the members of RLE who share the same floor in Building 36 and did study groups and class projects together, the members of a poker club who provided fun and entertainment every week, and my old friends Jae-Byum, Keun Sup, and Donghyun who I shared every joy and adversity with.

Most importantly, I thank my family. My parents and sister have supported me

not just for the thesis but for my entire life. Their trust has given me strength every step of my journey. And I thank Hyojin for being my wife.

- This research was supported by the National Science Foundation under Grant No. 1101147 and the Korea Foundation for Advanced Studies.

Contents

1	Introduction	15
1.1	Interactions among Multiple Agents	17
1.2	Outline and Contributions	19
2	Background	23
2.1	Bayesian Hypothesis Testing	23
2.1.1	Binary Hypothesis Testing	24
2.1.2	Maximum a Posteriori and Maximum Likelihood Decision Rules	26
2.1.3	The Operating Characteristic of the Likelihood Ratio Test . . .	26
2.2	Distributed Detection and Data Fusion	29
2.2.1	Optimal Fusion Rule and Local Decision Rule	29
2.2.2	Symmetric Fusion Rule	31
2.3	Social Learning	34
2.3.1	Social Learning and Belief Update	35
2.3.2	Herding Behavior	36
3	Human Perception of Prior Probability	41
3.1	Studies on Human Decision Making	42
3.2	Estimation of Prior Probabilities	45
4	Distributed Detection with Symmetric Fusion	49
4.1	An Imperfect Agent	50
4.1.1	The Number of Observations	50
4.1.2	The Number of Agents	52

4.2	Team of Imperfect Agents	54
4.2.1	A Team of Agents Who Observe Conditionally IID Signals . .	55
4.2.2	A Team of Agents Who Observe the Same Signal	60
4.3	Experts and Novices in a Team	62
4.4	Conclusion	65
4.A	Proof of Theorem 4.1	67
5	Sequential Detection with Social Learning	71
5.1	Social Learning of Imperfect Agents	72
5.1.1	Alexis, the First Agent	74
5.1.2	Britta, the Second Agent	74
5.1.3	Carol, the Third Agent	76
5.1.4	Norah, the N th Agent	77
5.2	Optimal Initial Belief	79
5.2.1	An Open-Minded Advisor	81
5.2.2	An Informative Public Signal	82
5.3	Conclusion	87
5.A	Proof of Theorem 5.2	88
5.B	Proof of Theorem 5.3	90
6	Distributed Detection with Social Learning and Symmetric Fusion	95
6.1	Conditionally IID Private Signals and Complete Public Signals	96
6.1.1	Two-Agent Case	97
6.1.2	N -Agent Case	99
6.1.3	N -Imperfect-Agent Case	105
6.2	Conditionally IID Private Signals and Incomplete Public Signals . . .	105
6.3	Agents with Different Likelihoods	107
6.3.1	AND and OR Fusion Rules	108
6.3.2	Other Fusion Rules	110
6.4	Experimental Investigation of Human Social Learning in Team Deci- sion Making	113

6.4.1	Experiment A: Setup	113
6.4.2	Experiment A: Results	115
6.4.3	Experiment B: Setting and Results	118
6.5	Conclusion	120
6.A	Results of Experiment A	122
6.B	Results of xperiment B	126
7	Conclusion	131
7.1	Recapitulation	132
7.2	Future Directions	134

List of Figures

2-1	A receiver operating characteristic (ROC) curve	27
2-2	Concavity of the ROC curve	28
3-1	An extensive hypothesis testing model	46
4-1	An agent with multiple observations	50
4-2	The minimum number of observations for an imperfect agent to outperform a perfect agent	52
4-3	Multiple agents forming a team	53
4-4	The minimum number of imperfect agents to outperform a perfect agent	54
4-5	Comparison of the Bayes risk of a team of individual agents to that of a team of diverse agents	56
4-6	A contour plot of the Bayes risk of a two-agent team with the OR fusion rule	59
4-7	Decision making when agents have the same observations	60
4-8	The Bayes risk of a team consisting of an expert and a novice	65
4-9	The expert's and the novice's beliefs that minimize the team Bayes risk	66
5-1	A sequential decision making model with N agents	72
5-2	Diana's belief update	78
5-3	Britta's Bayes risk depending on Alexis's and Britta's beliefs	81
5-4	The trend of Alexis's and Britta's optimal beliefs that minimize Britta's Bayes risk	81

5-5	The trend of Alexis's, Britta's, and Carol's optimal beliefs that minimize Carol's Bayes risk	83
5-6	Mutual information $I(H; \widehat{H}_1)$ when Alexis is open-minded	84
5-7	Alexis's belief that maximizes her selfless utility	85
5-8	Increase of Alexis's Bayes risk and decrease of Britta's Bayes risk due to Alexis's being open-minded	87
5-9	The difference between the left-hand side and the right-hand side of (5.26)	91
5-10	An ROC curve that explains (5.34)	94
6-1	A distributed detection model with social learning and decision fusion	96
6-2	An $(N+1)$ -agent problem divided into two N -agent problems depending on Alexis's decision \widehat{H}_1	102
6-3	Illustration of the separate effects of the belief update and the fusion rule evolution.	104
6-4	Examples of partial public voting	106
6-5	Lower bounds of operating regions for a team of three agents with different noise variances	111
6-6	Lower bounds of operating regions for a team of four agents with different noise variances	112
6-7	The likelihoods of a private signal in Experiment A	114
6-8	Optimal decision-making thresholds and empirical thresholds for various public signals in Experiment A	116
6-9	Comparison of decision thresholds in Version I and Version II in Experiment A	117
6-10	Optimal decision-making rules and empirical rules for various public signals in Experiment B	120

List of Tables

4.1	Variance of the median of N iid Gaussian random variables	50
6.1	Subjects' Decision Thresholds When They Draw Nine Balls	122
6.2	Subjects' Decision Thresholds When They Draw One Ball	127

Introduction

A perfect decision-making system processes information with full knowledge of its statistical properties to make the best possible decisions. The processing includes observing signals, computing sufficient statistics, comparing the outcomes of decisions, and maximizing the payoffs. However, some systems are not perfect because of limitations on information, memory size, computing power, or energy. This thesis deals with decision systems that have cognitive limitations and consequently incorrectly perceive the prior probabilities of the given alternatives. The best performances of these imperfect decision systems are found and compared to those of perfect decision systems. The thesis focuses on decision-making problems performed by multiple agents rather than by a single agent because a group of multiple agents has more flexibility to cope with the incorrect knowledge of prior probabilities. The effect of incorrect prior probabilities is discussed in three fundamental forms of decision fusion—parallel decision making with a symmetric fusion rule, sequential decision making, and sequential decision making with a symmetric fusion rule.

Detection theory has been developed to improve techniques that distinguish information from noisy signals. The primary objective of a detection rule is to decrease the possibility or the expected cost of incorrect detection of target states. The detection rule consists of a set of thresholds to classify the noisy signals into one of several states. It is important for a decision system to determine the thresholds that meet a desired criterion.

The design of thresholds requires some statistics about the states and signals. It requires the likelihood function of the signals conditioned on each state for the

maximum likelihood criterion. It requires the likelihood functions of the signals and the prior probabilities of states for the minimum error probability criterion. Both statistics are also required for the minimum Bayes risk criterion.

The minimum Bayes risk criterion evaluates detection performance with respect to the expected decision-making cost, which is called Bayes risk. Each type of error may induce a different cost in different situations. For example, in a medical diagnosis problem where the hypothesis is that a patient has a particular disease, a missed detection of the disease will cause much more damage than a false alarm. Conversely, in a criminal trial where the hypothesis is that the defendant is guilty, a false alarm is considered more seriously than a missed detection. A Bayesian decision-making agent knows proper costs of errors for each decision-making task and behaves to minimize the expected cost. Rational human agents are known to approximately follow the minimum Bayes risk criterion [1–4].

However, the prior probabilities are not easily known in practice, especially when it comes to human decision-making tasks. What would be the prior probability that a certain person commits a crime? What would be the prior probability that a job applicant will make a huge contribution to our company? People do not know and cannot even calculate these prior probabilities but still need them in order to reach a more accurate verdict or to hire the best applicant.

Psychological studies argue that people use the information that they can observe to perceive the prior probabilities. For example, a juror measures from a criminal defendant’s appearance the degree to which the defendant looks intelligent, friendly, attractive, and violent [5, 6]. These measures are correlated with the verdicts—guilt or innocence. We believe that the measures are mapped to a value that represent the juror’s estimation of the prior probability. Even though the value is not equal to the true prior probability, it is what the juror uses in his or her decision making. Since, strictly speaking, the estimation is not the same as the prior probability, it is referred to as *perceived belief* or simply *belief* throughout the thesis. In the language of economics, humans know their payoff functions, which is negative Bayes risk, and rational human beings behave to maximize their payoffs. However, from our per-

spective, what they know is different from what the functions really are. There is a mismatch between people’s perceptions and the reality.

The mismatch will prevent humans from performing optimal decision making. While the prior probability is one of the basic elements of detection, the effect of accuracy of the prior probability has not received a great deal of attention. It is not a solution to encourage people to more accurately estimate the prior probability. This problem should be dealt from a decision-making system level. In other words, we admit that people are imperfect decision makers and try to design a decision-making system that relieves or supplements the imperfectness. As a first step, this thesis pursues to understand the effect of the limitation on the perception of prior probability, to analyze decision-making performance of multiple imperfect agents in various topologies, and to optimize their decision strategies under the limitation from a perspective of costly rationality [7].

Section 1.1 explains the motivation to have an interest in a group of decision-makers and the diversity in it. Section 1.2 summarizes the main contributions of the thesis.

■ 1.1 Interactions among Multiple Agents

Most of this thesis is focused on decision-making agents, with beliefs that are not necessarily correct, forming a team and interacting. The natural effect of the interaction is to influence each others beliefs, and the ramifications for the team’s performance are intricate.

Previously, we have considered categorically-thinking agents. Under this model, initially proposed in [8], the agents can only memorize and process at most K different prior probability values. This limitation was motivated by a scenario where an agent needs to make decisions for a lot of objects with different prior probabilities. It is impractical and inefficient for the agent to perform accurate likelihood ratio tests for each decision making task because computational resources are required proportional to the number of decision-making subjects. Thus, the agent first classifies similar

objects to the same category and design only one decision rule for the category. The agent, who can deal with at most K categories, can process infinitely many subjects with the corresponding K decision rules. In fact, this is a natural method that human decision makers use [9].

The classification degrades each agent's decision making whereas it reduces computational complexity. For each subject of the hypothesis testing, the applied decision rule is not optimally designed for the prior probability of the individual subject; it is designed for the category that the subject belongs to. Therefore, category-wise optimal decision rules were found instead of subject-wise optimal decision rules in [8]. They model the prior probability as a random variable with a given distribution and investigate conditions of the optimal quantizers for the prior probability when the optimality criterion is the Bayes risk averaged over the prior probability. Then the average performance of decision making can be maximized.

One result for multiple categorically-thinking agents is especially inspiring for this thesis. Suppose the agents have conditionally independent and identically distributed (iid) observations, conditioned on the true state. While their team will get a benefit over any single agent alone from the multiplicity of agents even if they use the same categorization (quantization of prior probabilities), the benefit can be larger when the categorizations are different [10–12]. In effect, the particular cognitive limitation of categorical thinking is reduced through teamwork. Furthermore, the benefit from diversity is present even under limited coordination [13].

Inspired by this observation, our main concern is the effect of diverse recognition of prior probabilities on the decision making and the optimal form of diversity for the best decision making. This thesis discusses three topologies of the team of agents: a distributed detection and data fusion model, a sequential detection model, and a model combining these two. Agents behave differently in each model. For example, agents in the sequential network perform social learning while those in the parallel network perform decision fusion. This thesis finds the optimal decision strategies and the optimal form of decision aggregation in each model.

■ 1.2 Outline and Contributions

The central goal of the thesis is to provide theoretical results on the optimal binary decision-making strategies with incorrect prior probabilities. Instead of considering a trivial case with a single imperfect agent, the thesis focuses on multiple-agent cases under the following three topologies:

- **Distributed detection:** Agents individually observe signals and perform hypothesis testing but share their decision-making cost. Their detection results are sent to a fusion center and converted to a global decision according to a symmetric L -out-of- N fusion rule. The concern is to minimize the Bayes risk of the global decision.
- **Sequential detection:** Agents individually observe signals and sequentially perform hypothesis testing. Their decisions are available to other agents so any agent before making a decision can do social learning. The concern is to minimize the Bayes risk of the last-acting agent.
- **Distributed detection with social learning:** The model is similar to the distributed detection model except that the local decisions are available to all agents who make decisions sequentially. Thus the agents can do social learning like in the sequential detection model. The concern is to minimize the Bayes risk of the global decision.

Below is the outline of the thesis.

Chapter 2: Background

The background chapter reviews relevant results in detection theory and social learning. In particular, hypothesis testing problems are described and Bayes risk criterion is defined. Properties of the operating characteristic of likelihood ratio test are explained. The optimal fusion rule and local decision rules are obtained in a distributed detection problem. Finally, social learning is interpreted in terms of belief update and incorrect herding behavior is explained.

Chapter 3: Human Perception of Prior Probability

The chapter delivers our motivation of questioning human ability to perceive prior probability. Reviewing through a flow of several psychological studies from general behaviors of human decision makers to specific observations related to their perceiving prior probability, we develop a model of human perception. The Bayes-optimal estimation of prior probability is obtained, which justifies our imperfect agent model.

Chapter 4: Distributed Detection with Symmetric Fusion

The binary distributed detection problem is discussed with imperfect agents. Specifically, the fusion rule is restricted to be symmetric in a form of L -out-of- N rule that chooses 1 if L or more agents choose 1 and 0 otherwise. The symmetric rules are not always optimal but they are widely used in human decision-making scenarios.

The distributed detection of imperfect agents is outperformed by that of perfect agents when the agents observe conditionally iid signals. It is not possible to make a team of imperfect agents outperform a team of perfect agents, but it is possible to make the imperfect team more robust to the change of their beliefs by achieving diversity.

On the other hand, if the signals are not identically distributed, decision making of the perfect agents is always suboptimal. They cannot properly optimize their decision rules because they do not know the likelihood functions of other agents' private signals. However, since decision rules of imperfect agents depend on their beliefs as well, they can have either better or worse decision rules. It can happen that their wrong beliefs accidentally lead to the optimal decision rules.

Chapter 5: Sequential Detection with Social Learning

Agents can observe others' decisions as *public signals* in this sequential detection model. The public signals give different information so the agents can learn from them to improve their decision rules. This behavior, which is called *social learning* or *observational learning*, is the main interest of Chapter 5.

When the agents have different incorrect beliefs, their social learning becomes in-

complete. Unlike the intuition that the incomplete social learning and incorrect belief should cause increase of Bayes risk, however, some incorrect beliefs reduced the Bayes risk even below that of perfect agents. It is proven that such incorrect beliefs have a systematic pattern. The results are summarized in the statement, “Earlier-acting agents should be open-minded,” which implies that they should overweight small probabilities and underweight high probabilities. The open-mindedness is related to being informative to later-acting agents.

Chapter 6: Distributed Detection with Social Learning and Symmetric Fusion

Social learning generally helps individual decision makers reduce their Bayes risks. It is examined whether social learning also reduces Bayes risk of a team in a distributed detection model that allows agents to observe others’ decisions and do social learning. It turns out that social learning does not lead to the improvement of team performance in some cases especially when the agents observe conditionally iid private signals. It is proven that it is the optimal behavior to ignore public signals and not to do social learning in that case. In other cases when the agents observe private signals not identically distributed, social learning generally improves team decision-making performance and the order of decision making matters to the team performance.

Human behaviors are tested in the case when they are given conditionally iid private signals and some public signals in team decision-making tasks by experiments conducted on Amazon Mechanical Turk. The experiments reveal that humans use public signals even when the rational behavior is to ignore them.

Chapter 7: Conclusion

The last chapter concludes the thesis with recapitulation of the main results and their practical impacts. Besides, future directions of the study of imperfect agents and social learning are proposed.

Bibliographical Note

Parts of Chapter 5 appear in the paper:

- J. B. Rhim and V. K. Goyal, “Social Teaching: Being Informative vs. Being Right in Sequential Decision Making,” in *Proceedings of the IEEE International Symposium on Information Theory*, pp. 2602–2606, July 2013.

Parts of Chapter 6 appear in the papers:

- J. B. Rhim and V. K. Goyal, “Social Learning in Team Decision Making,” to appear in *Proceedings of Collective Intelligence Conference*, June 2014.
- J. B. Rhim and V. K. Goyal, “Keep Ballots Secret: On the Futility of Social Learning in Decision Making by Voting,” in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 4231–4235, May 2013.

and in the journal manuscript:

- J. B. Rhim and V. K. Goyal, “Distributed Hypothesis Testing with Social Learning and Symmetric Fusion,” arXiv:1404.0964 [cs.IT].

Background

The rational decision-making strategy is to exploit available information and maximize payoffs. This background chapter reviews basic concepts in detection theory that have been developed to describe such rational behavior. First, Bayesian hypothesis testing problems of a single agent and several performance criteria are presented. Then distributed detection problems with multiple agents are reviewed. Finally, social learning among agents is discussed, including an explanation of the possibility of incorrect herding.

Throughout this thesis, $f(\cdot)$ denotes distributions of continuous random variables and $p(\cdot)$ distributions of discrete random variables.

■ 2.1 Bayesian Hypothesis Testing

There are many applications of decision making with a set of observations. Customers purchase one of several products, juries reach verdicts, physicians make diagnoses, and receivers recover signals. The decision-making agents generally observe noisy signals. They are trying to use appropriate decision process to make correct or best decisions as much as possible from the noisy observations.

Decision making is mathematically abstracted as hypothesis testing. There are M hypotheses or states $\mathcal{H} = \{0, 1, \dots, M - 1\}$ from which an agent can choose. Only one hypothesis is true and the model of the observed signal $Y \in \mathcal{Y}$, which is represented as a random variable or random vector, is determined by the true hypothesis.

In the Bayesian approach, a complete characterization of knowledge about H upon observing $Y = y$ is described by posterior distributions $p_{H|Y}(h|y)$, for $h =$

$0, 1, \dots, M - 1$. Bayes rule is used to compute the posterior distributions. The computation requires a prior distribution of H

$$p_H(h), \quad h = 0, 1, \dots, M - 1,$$

and characterizations of the observed data under each hypothesis

$$f_{y|H}(y|h), \quad h = 0, 1, \dots, M - 1.$$

■ 2.1.1 Binary Hypothesis Testing

Choosing between two alternatives is the simplest form of hypothesis testing but is very common. In the binary case, two hypotheses $H = 0$ and $H = 1$ are considered and their prior probabilities are denoted by $p_0 = p_H(0)$ and $p_1 = p_H(1) = 1 - p_0$. The information Y is described by one of two likelihood functions $f_{Y|H}(y|0)$ and $f_{Y|H}(y|1)$. The solution to a hypothesis testing is given by a decision rule $\widehat{H} : \mathcal{Y} \mapsto \{0, 1\}$.

In the Bayesian approach, the expected value of decision-making cost is the criterion for a good decision rule. Specifically, we use a cost function $C(j, i) = c_{ji}$ to denote the cost of deciding $\widehat{H} = j$ when the correct hypothesis is $H = i$. The expected cost is

$$R = \mathbb{E}[C(j, i)] = \sum_{i=0}^1 \sum_{j=0}^1 c_{ji} \mathbb{P}\{\widehat{H} = j, H = i\},$$

which is called Bayes risk. The optimal decision rule takes the form

$$\begin{aligned} \widehat{H}(\cdot) &= \arg \min R \\ &= \arg \min \sum_{i=0}^1 \sum_{j=0}^1 c_{ji} \int_{\mathcal{Y}_j} f_{Y,H}(y, i) dy, \end{aligned}$$

where $\mathcal{Y}_j = \{y | \widehat{H}(y) = j\}$ denotes the region of the signal space \mathcal{Y} that corresponds to decision $\widehat{H} = j$; \mathcal{Y}_0 and \mathcal{Y}_1 form a partition of \mathcal{Y} .

Consider an arbitrary but fixed decision rule $\widehat{H}(\cdot)$. The Bayes risk can be expanded

to the form

$$R = \int_{\mathcal{Y}_0} (c_{00}p_0 f_{Y|H}(y|0) + c_{01}(1-p_0) f_{Y|1}(y|1)) dy + \int_{\mathcal{Y}_1} (c_{10}p_0 f_{Y|H}(y|0) + c_{11}(1-p_0) f_{Y|1}(y|1)) dy.$$

For $\widehat{H}(\cdot)$ to minimize the Bayes risk, the expression for R can be minimized pointwise over y by comparing the possible contributions and choosing whichever has a smaller value:

$$c_{00}p_0 f_{Y|H}(y|0) + c_{01}(1-p_0) f_{Y|1}(y|1) \underset{\widehat{H}(y)=0}{\overset{\widehat{H}(y)=1}{\gtrless}} c_{10}p_0 f_{Y|H}(y|0) + c_{11}p_1 f_{Y|1}(y|1).$$

Collecting terms and rewriting as a ratio gives a form of the likelihood ratio test (LRT)

$$\frac{f_{Y|H}(y|1)}{f_{Y|H}(y|0)} \underset{\widehat{H}(y)=0}{\overset{\widehat{H}(y)=1}{\gtrless}} \frac{(c_{10} - c_{00})p_0}{(c_{01} - c_{11})(1-p_0)} \triangleq \eta. \quad (2.1)$$

The optimal strategy is a deterministic decision rule. Bayesian agents perform the LRT to make optimal decisions.

The thesis constrains the likelihood functions so that their ratio $f_{Y|H}(y|1)/f_{Y|H}(y|0)$ is monotonically increasing in y . It simplifies the LRT (2.1) to the comparison of the signal to a threshold:

$$y \underset{\widehat{H}(y)=0}{\overset{\widehat{H}(y)=1}{\gtrless}} \lambda. \quad (2.2)$$

In other words, $\mathcal{Y}_0 = (-\infty, \lambda)$ and $\mathcal{Y}_1 = [\lambda, \infty)$. For example, if $Y = H + W$, where W is a Gaussian random variable with zero mean and variance σ^2 , the LRT (2.1) can be expressed as

$$y \underset{\widehat{H}(y)=0}{\overset{\widehat{H}(y)=1}{\gtrless}} \frac{1}{2} + \sigma^2 \ln \frac{(c_{10} - c_{00})p_0}{(c_{01} - c_{11})(1-p_0)}.$$

The threshold λ is called *decision threshold* throughout the thesis.

■ 2.1.2 Maximum a Posteriori and Maximum Likelihood Decision Rules

The likelihood ratio test (2.1) is the optimal decision rule for any cost function $C(\cdot, \cdot)$.

In a special case when

$$c_{ji} = \begin{cases} 0, & i = j, \\ 1, & i \neq j, \end{cases} \quad (2.3)$$

the Bayes cost is equal to the probability of errors $\mathbb{P}\{\widehat{H} = 1, H = 0\} + \mathbb{P}\{\widehat{H} = 0, H = 1\}$.

The decision rule (2.1) is expressed in terms of the a posteriori probabilities:

$$p_{H|Y}(1|y) \underset{\widehat{H}(y)=0}{\overset{\widehat{H}(y)=1}{\gtrless}} p_{H|Y}(0|y).$$

This rule, which chooses the hypothesis whose a posterior probability is larger, is called the maximum a posteriori (MAP) decision rule.

Furthermore, the decision rule gets simpler if the hypotheses are equally likely:

$$f_{Y|H}(y|1) \underset{\widehat{H}(y)=0}{\overset{\widehat{H}(y)=1}{\gtrless}} f_{Y|H}(y|0).$$

This rule chooses the hypothesis for which the corresponding likelihood function is larger. It is called the maximum likelihood (ML) decision rule.

For simplicity, it is assumed that $c_{00} = c_{11} = 0$ throughout the thesis. However, the costs of errors c_{10} and c_{01} are not restricted to be 1; the LRT (2.1) will be used.

■ 2.1.3 The Operating Characteristic of the Likelihood Ratio Test

The performance of any decision rule $\widehat{H}(\cdot)$ is specified in terms of two quantities

$$\begin{aligned} P_e^I &= \mathbb{P}\{\widehat{H}(Y) = 1 | H = 0\} = \int_{\mathcal{Y}_1} f_{Y|H}(y|0) dy, \\ P_e^{II} &= \mathbb{P}\{\widehat{H}(Y) = 0 | H = 1\} = \int_{\mathcal{Y}_0} f_{Y|H}(y|1) dy, \end{aligned} \quad (2.4)$$

where P_e^I is called false alarm (or Type I error) probability and P_e^{II} missed detection (or Type II error) probability. Another set of notations $P_D = \mathbb{P}\{\widehat{H}(Y) = 1 | H = 1\} = 1 - P_e^{II}$

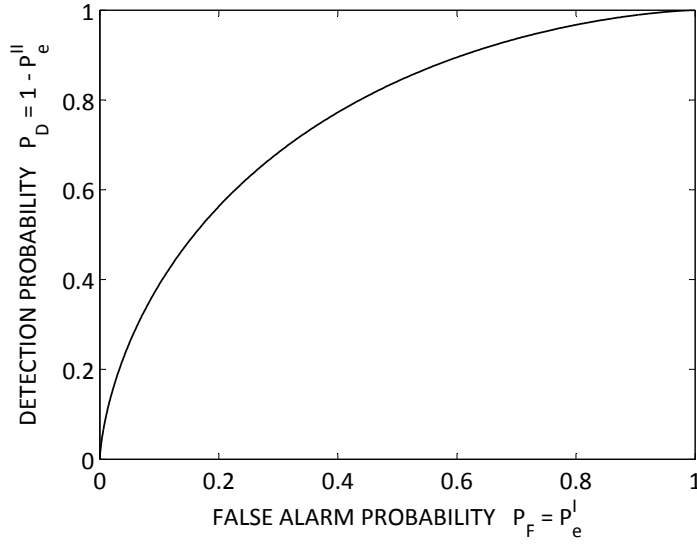


Figure 2-1. Receiver operating characteristic curve of the likelihood ratio test for a detection problem with additive Gaussian noise. The noise has zero mean and unit variance.

and $P_F = \mathbb{P}\{\widehat{H}(Y) = 1 | H = 0\} = P_e^I$ are also widely used but P_e^I and P_e^{II} will be dominantly used throughout the thesis.

A good decision rule is one with small P_e^I and P_e^{II} . However, these are competing objectives. Increasing the threshold η in (2.1), or equivalently λ in (2.2), causes increase of the length of \mathcal{Y}_0 and decrease of the length of \mathcal{Y}_1 . Accordingly, P_e^{II} is increased and P_e^I is decreased. Thus, choosing a threshold λ involves making an acceptable tradeoff between P_e^I and P_e^{II} .

As λ is varied from $-\infty$ to ∞ , a curve for (P_D, P_F) is traced out in the space between $[0, 0]$ and $[1, 1]$ as depicted in Figure 2-1. This curve, which is referred to as the receiver operating characteristic (ROC) curve, characterizes the likelihood ratio test.

The ROC curve of a likelihood ratio test has the following properties:

- (a) It always contains the two points $(0, 0)$ and $(1, 1)$.
- (b) It is monotonic.
- (c) $P_D \geq P_F$.
- (d) It is concave.




Figure 2-2. A non-concave ROC curve can be turned into a concave curve by using a randomized test. Take the two end points $(P_F(\lambda_1), P_D(\lambda_1))$ and $(P_F(\lambda_2), P_D(\lambda_2))$ of the interval where the ROC curve is not concave. A randomized test that uses the threshold λ_1 with probability p and λ_2 with probability $1 - p$, for $p \in [0, 1]$, will achieve the dashed line.

Property (a) is obvious; $(P_F, P_D) \rightarrow (0, 0)$ as $\lambda \rightarrow -\infty$ and $(P_F, P_D) \rightarrow (1, 1)$ as $\lambda \rightarrow \infty$. Property (b) comes immediately from the structure of the likelihood ratio test. Let $P_F(\lambda)$ and $P_D(\lambda)$ respectively denote the false alarm and detection probability of the test (2.1) associated with a threshold λ . For any λ_1 and λ_2 such that $\lambda_2 > \lambda_1$, $P_F(\lambda_1) \geq P_F(\lambda_2)$ and $P_D(\lambda_1) \geq P_D(\lambda_2)$. Therefore, the slope of the ROC curve is nonnegative:

$$\frac{P_D(\lambda_2) - P_D(\lambda_1)}{P_F(\lambda_2) - P_F(\lambda_1)} \geq 0.$$

Property (c) is achieved by flipping decisions if $P_D < P_F$. Then the new false alarm probability becomes $1 - P_F$ and detection probability becomes $1 - P_D$. The reversed test will have better performance because its operating characteristic is above the line $P_F = P_D$.

Property (d) is true if the likelihood ratio is monotonically increasing. The slope of the ROC curve at a point $(P_F(\lambda), P_D(\lambda))$ is

$$\frac{dP_D(\lambda)}{dP_F(\lambda)} = -\frac{dP_e^{\text{II}}(\lambda)}{dP_e^{\text{I}}(\lambda)} = -\frac{dP_e^{\text{II}}/d\lambda}{dP_e^{\text{I}}/d\lambda} = \frac{f_{Y|H}(\lambda|1)}{f_{Y|H}(\lambda|0)},$$

which is nonnegative and monotonically increasing in λ by the given condition. Since $P_F(\lambda)$ is a decreasing function of λ , the slope dP_D/dP_F is monotonically decreasing in P_F . Therefore, the ROC curve is concave. If the likelihood ratio is not monotonically

increasing, the LRT may be replaced with a randomized test to replace a non-concave ROC curve with its convex hull (see Figure 2-2).

■ 2.2 Distributed Detection and Data Fusion

Decision making in large-scale systems may consist of multiple decision makers to improve system performance. A distributed detection system consists of multiple agents that observe signals and make decisions in parallel. The decisions are sent to a fusion center that makes a global decision [14, 15].

■ 2.2.1 Optimal Fusion Rule and Local Decision Rule

The data fusion is also a binary hypothesis testing problem in which observations are the local decisions. The optimal fusion rule for N agents, $\widehat{H} : \{0, 1\}^N \mapsto \{0, 1\}$, is given by the following likelihood ratio test

$$\frac{p_{\widehat{H}_1, \dots, \widehat{H}_N | H}(\widehat{h}_1, \dots, \widehat{h}_N | 1)}{p_{\widehat{H}_1, \dots, \widehat{H}_N | H}(\widehat{h}_1, \dots, \widehat{h}_N | 0)} \underset{\widehat{H}=0}{\overset{\widehat{H}=1}{\gtrless}} \frac{(c_{10} - c_{00})p_0}{(c_{01} - c_{11})(1 - p_0)}.$$

If the agents observe conditionally independent signals, their local decisions are conditionally independent as well. The corresponding log-likelihood ratio test is

$$\sum_{n=1}^N \log \frac{p_{\widehat{H}_n | H}(\widehat{h}_n | 1)}{p_{\widehat{H}_n | H}(\widehat{h}_n | 0)} \underset{\widehat{H}=0}{\overset{\widehat{H}=1}{\gtrless}} \log \frac{(c_{10} - c_{00})p_0}{(c_{01} - c_{11})(1 - p_0)}.$$

Since

$$\frac{p_{\widehat{H}_n | H}(0 | 1)}{p_{\widehat{H}_n | H}(0 | 0)} = \frac{P_{e,n}^{\text{II}}}{1 - P_{e,n}^{\text{I}}} \quad \text{and} \quad \frac{p_{\widehat{H}_n | H}(1 | 1)}{p_{\widehat{H}_n | H}(1 | 0)} = \frac{1 - P_{e,n}^{\text{II}}}{P_{e,n}^{\text{I}}},$$

the data fusion rule is expressed as¹

$$\sum_{n=1}^N w_n (2\widehat{h}_n - 1) \underset{\widehat{H}=0}{\overset{\widehat{H}=1}{\gtrless}} \eta, \quad (2.5)$$

¹This expression is different from the fusion rule in [14] because $H \in \{-1, +1\}$ in the paper but $H \in \{0, 1\}$ in this thesis.

where

$$w_n = \begin{cases} \log \frac{1 - P_{e,n}^I}{P_{e,n}^{II}}, & \text{if } \widehat{h}_n = 0, \\ \log \frac{1 - P_{e,n}^{II}}{P_{e,n}^I}, & \text{if } \widehat{h}_n = 1, \end{cases} \quad \text{and} \quad \eta = \log \frac{(c_{10} - c_{00})p_0}{(c_{01} - c_{11})(1 - p_0)}.$$

The optimal fusion rule is comparison of weighted sum of local decisions to the threshold η .

Deriving the optimal local decision rule of Agent n , $\widehat{H}_n : \mathcal{Y} \mapsto \{0, 1\}$, assumes that the fusion rule and all the other local decision rules have been already designed and are remaining fixed. The Bayes risk is expressed in terms of decisions of Agent n as

$$\begin{aligned} R &= (c_{10} - c_{00})p_0 p_{\widehat{H}|H}(1|0) + (c_{11} - c_{01})(1 - p_0) p_{\widehat{H}|H}(1|1) + (c_{00}p_0 + c_{01}(1 - p_0)) \\ &= \sum_{\widehat{\mathbf{h}}_{-n}} C_F (p_{\widehat{H}, \widehat{\mathbf{H}}_{-n}, \widehat{H}_n|H}(1, \widehat{\mathbf{h}}_{-n}, 0|0) + p_{\widehat{H}, \widehat{\mathbf{H}}_{-n}, \widehat{H}_n|H}(1, \widehat{\mathbf{h}}_{-n}, 1|0)) \\ &\quad - \sum_{\widehat{\mathbf{h}}_{-n}} C_D (p_{\widehat{H}, \widehat{\mathbf{H}}_{-n}, \widehat{h}_n|H}(1, \widehat{\mathbf{h}}_{-n}, 0|1) + p_{\widehat{H}, \widehat{\mathbf{H}}_{-n}, \widehat{h}_n|H}(1, \widehat{\mathbf{h}}_{-n}, 1|1)) + C, \end{aligned}$$

where $C_F \triangleq (c_{10} - c_{00})p_0$, $C_D \triangleq (c_{01} - c_{11})(1 - p_0)$, $C \triangleq c_{00}p_0 + c_{01}(1 - p_0)$, and

$$\widehat{\mathbf{H}}_{-n} \triangleq \{\widehat{H}_1, \dots, \widehat{H}_{n-1}, \widehat{H}_{n+1}, \dots, \widehat{H}_N\}.$$

Since $p_{\widehat{H}|\widehat{\mathbf{H}}_{-n}, \widehat{H}_n, H}(\widehat{h}|\widehat{\mathbf{h}}_{-n}, \widehat{h}_n, h) = p_{\widehat{H}|\widehat{\mathbf{H}}_{-n}, \widehat{H}_n}(\widehat{h}|\widehat{\mathbf{h}}_{-n}, \widehat{h}_n)$,

$$\begin{aligned} R &= \sum_{\widehat{\mathbf{h}}_{-n}} \left\{ [p_{\widehat{H}|\widehat{\mathbf{H}}_{-n}, \widehat{H}_n}(1|\widehat{\mathbf{h}}_{-n}, 1) - p_{\widehat{H}|\widehat{\mathbf{H}}_{-n}, \widehat{H}_n}(1|\widehat{\mathbf{h}}_{-n}, 0)] \right. \\ &\quad \times [C_F p_{\widehat{\mathbf{H}}_{-n}, \widehat{H}_n|H}(\widehat{\mathbf{h}}_{-n}, 1|0) - C_D p_{\widehat{\mathbf{H}}_{-n}, \widehat{H}_n|H}(\widehat{\mathbf{h}}_{-n}, 1|1)] \left. \right\} \\ &\quad + \sum_{\widehat{\mathbf{h}}_{-n}} p_{\widehat{H}|\widehat{\mathbf{H}}_{-n}, \widehat{H}_n}(1|\widehat{\mathbf{h}}_{-n}, 0) [C_F p_{\widehat{\mathbf{H}}_{-n}|H}(\widehat{\mathbf{h}}_{-n}|0) - C_D p_{\widehat{\mathbf{H}}_{-n}|H}(\widehat{\mathbf{h}}_{-n}|1)] + C \end{aligned}$$

The terms corresponding to all agents other than Agent n are constant. Minimizing the Bayes risk is equivalent to minimizing

$$\sum_{\widehat{\mathbf{h}}_{-n}} A(\widehat{\mathbf{h}}_{-n}) \left[C_F \prod_{m \neq n} p_{\widehat{H}_m|H}(\widehat{h}_m|0) p_{\widehat{H}_n|H}(1|0) - C_D \prod_{m \neq n} p_{\widehat{H}_m|H}(\widehat{h}_m|1) p_{\widehat{H}_n|H}(1|1) \right],$$

where

$$A(\widehat{\mathbf{h}}_{-n}) \triangleq p_{\widehat{H}|\widehat{\mathbf{h}}_{-n},\widehat{H}_n}(1|\widehat{\mathbf{h}}_{-n},1) - p_{\widehat{H}|\widehat{\mathbf{h}}_{-n},\widehat{H}_n}(1|\widehat{\mathbf{h}}_{-n},0). \quad (2.6)$$

The Bayes risk is minimized by the local decision rule that decides 1 if

$$\sum_{\widehat{\mathbf{h}}_{-n}} A(\widehat{\mathbf{h}}_{-n}) C_F \prod_{k \neq n} p_{\widehat{H}_k|H}(\widehat{h}_m|0) f_{Y_n|H}(y_n|0) \leq \sum_{\widehat{\mathbf{h}}_{-n}} A(\widehat{\mathbf{h}}_{-n}) C_D \prod_{m \neq n} p_{\widehat{H}_m|H}(\widehat{h}_m|1) f_{Y_n|H}(y_n|1)$$

and decides 0 otherwise. Thus the optimal local decision rule is

$$\frac{f_{Y_n|H}(y_n|1)}{f_{Y_n|H}(y_n|0)} \underset{\widehat{H}_n(y_n)=0}{\overset{\widehat{H}_n(y_n)=1}{\gtrless}} \frac{C_F \sum_{\widehat{\mathbf{h}}_{-n}} A(\widehat{\mathbf{h}}_{-n}) \prod_{k \neq n} p_{\widehat{H}_k|H}(\widehat{h}_m|0)}{C_D \sum_{\widehat{\mathbf{h}}_{-n}} A(\widehat{\mathbf{h}}_{-n}) \prod_{m \neq n} p_{\widehat{H}_m|H}(\widehat{h}_m|1)}. \quad (2.7)$$

The optimal local decision rules and fusion rule are the solution to (2.5) and (2.7), a total $N + 1$ equations.

■ 2.2.2 Symmetric Fusion Rule

Local decisions are often fused by a symmetric L -out-of- N rule. The global decision is 1 if L or more agents decide 1 and is 0 otherwise, i.e.,

$$\sum_{n=1}^N \widehat{H}_n \underset{\widehat{H}=0}{\overset{\widehat{H}=1}{\gtrless}} L.$$

Examples are the MAJORITY rule ($L = \lceil \frac{N+1}{2} \rceil$) and the OR rule ($L = 1$).

The optimal local decision rule is derived from (2.7). The L -out-of- N fusion rule yields

$$p_{\widehat{H}|\widehat{\mathbf{h}}_{-n},\widehat{H}_n}(1|\widehat{\mathbf{h}}_{-n},1) = \begin{cases} 1, & \text{if } \sum_{m \neq n} \widehat{h}_m \geq L - 1, \\ 0, & \text{if } \sum_{m \neq n} \widehat{h}_m < L - 1, \end{cases}$$

$$p_{\widehat{H}|\widehat{\mathbf{h}}_{-n},\widehat{H}_n}(1|\widehat{\mathbf{h}}_{-n},0) = \begin{cases} 1, & \text{if } \sum_{m \neq n} \widehat{h}_m \geq L, \\ 0, & \text{if } \sum_{m \neq n} \widehat{h}_m < L. \end{cases}$$

From (2.6),

$$A(\widehat{\mathbf{h}}_{-n}) = \begin{cases} 1, & \text{if } \sum_{m \neq n} \widehat{h}_m = L - 1, \\ 0, & \text{otherwise.} \end{cases} \quad (2.8)$$

Substituting (2.8) into (2.7) gives the optimal local decision rule of Agent n :

$$\frac{f_{Y_n|H}(y_n|1)}{f_{Y_n|H}(y_n|0)} \underset{\widehat{H}_n(y_n)=0}{\overset{\widehat{H}_n(y_n)=1}{\gtrless}} \frac{(c_{10} - c_{00})p_0 \sum_{\substack{I \subseteq [N] \setminus \{n\} \\ |I|=L-1}} \prod_{i \in I} p_{\widehat{H}_i|H}(1|0) \prod_{j \in [N] \setminus (I \cup \{n\})} p_{\widehat{H}_j|H}(0|0)}{(c_{01} - c_{11})(1 - p_0) \sum_{\substack{I \subseteq [N] \setminus \{n\} \\ |I|=L-1}} \prod_{i \in I} p_{\widehat{H}_i|H}(1|1) \prod_{j \in [N] \setminus (I \cup \{n\})} p_{\widehat{H}_j|H}(0|1)}, \quad (2.9)$$

where $[N]$ denotes the set $\{1, 2, \dots, N\}$.

In a special case when the signals Y_i are identically distributed conditioned on H , the agents are under identical conditions with respect to the quality of their information and the importance of their decisions. In this case, the constraint that the agents use identical decision rules causes little or no loss of performance [15, 16]. This constraint simplifies the problem of optimizing local decision rules because it significantly reduces the number of parameters.

The local decision rule of Agent n yields error probabilities like (2.4):

$$P_{e,n}^I = \mathbb{P}\{\widehat{H}_n = 1 | H = 0\},$$

$$P_{e,n}^{II} = \mathbb{P}\{\widehat{H}_n = 0 | H = 1\}.$$

It is easy to obtain closed forms of global error probabilities due to the symmetry of the fusion rule:

$$\begin{aligned} P_E^I &= \mathbb{P}\{\widehat{H} = 1 | H = 0\} = \mathbb{P}\{\sum_{m=1}^N \widehat{H}_m \geq L | H = 0\} \\ &= \sum_{m=L}^N \sum_{\substack{I \subseteq [N] \\ |I|=m}} \prod_{i \in I} P_{e,i}^I \prod_{j \in [N] \setminus I} (1 - P_{e,j}^I), \\ P_E^{II} &= \mathbb{P}\{\widehat{H} = 0 | H = 1\} = \mathbb{P}\{\sum_{m=1}^N \widehat{H}_m \leq L - 1 | H = 1\} \\ &= \sum_{m=N-L+1}^N \sum_{\substack{I \subseteq [N] \\ |I|=m}} \prod_{i \in I} P_{e,i}^{II} \prod_{j \in [N] \setminus I} (1 - P_{e,j}^{II}). \end{aligned}$$

For convenience, let us use the following notations to denote global error probability functions:

$$G_{L,N}^I(P_{e,1}^I, P_{e,2}^I, \dots, P_{e,N}^I) \triangleq P_E^I = \sum_{m=L}^N \sum_{\substack{I \subseteq [N] \\ |I|=m}} \prod_{i \in I} P_{e,i}^I \prod_{j \in [N] \setminus I} (1 - P_{e,j}^I), \quad (2.10)$$

$$G_{L,N}^{II}(P_{e,1}^{II}, P_{e,2}^{II}, \dots, P_{e,N}^{II}) \triangleq P_E^{II} = \sum_{m=N-L+1}^N \sum_{\substack{I \subseteq [N] \\ |I|=m}} \prod_{i \in I} P_{e,i}^{II} \prod_{j \in [N] \setminus I} (1 - P_{e,j}^{II}). \quad (2.11)$$

The global error probabilities are divided into two cases when Agent n is pivotal and when she is not:

$$\begin{aligned} G_{L,N}^I(P_{e,1}^I, P_{e,2}^I, \dots, P_{e,N}^I) &= \sum_{\substack{I \subseteq [N] \setminus \{n\} \\ |I|=L-1}} \prod_{i \in I} P_{e,i}^I \prod_{j \in [N] \setminus (I \cup \{n\})} (1 - P_{e,j}^I) P_{e,n}^I \\ &+ \sum_{m=L}^{N-1} \sum_{\substack{I \subseteq [N] \setminus \{n\} \\ |I|=m}} \prod_{i \in I} P_{e,i}^I \prod_{j \in [N] \setminus (I \cup \{n\})} (1 - P_{e,j}^I), \end{aligned}$$

where the first term is the probability that exactly $L-1$ agents among the $N-1$ agents other than Agent n cause false alarms and Agent n also causes a false alarm. The second term is the probability that L or more agents among the $N-1$ agents other than Agent n causes false alarms and the decision of Agent n is irrelevant. Likewise,

$$\begin{aligned} G_{L,N}^{II}(P_{e,1}^{II}, P_{e,2}^{II}, \dots, P_{e,N}^{II}) &= \sum_{\substack{I \subseteq [N] \setminus \{n\} \\ |I|=N-L}} \prod_{i \in I} P_{e,i}^{II} \prod_{j \in [N] \setminus (I \cup \{n\})} (1 - P_{e,j}^{II}) P_{e,n}^{II} \\ &+ \sum_{m=N-L+1}^{N-1} \sum_{\substack{I \subseteq [N] \setminus \{n\} \\ |I|=m}} \prod_{i \in I} P_{e,i}^{II} \prod_{j \in [N] \setminus (I \cup \{n\})} (1 - P_{e,j}^{II}). \end{aligned}$$

We use the notations $g_{L,N}^I$ and $g_{L,N}^{II}$ with the following definitions:

$$g_{L,N}^I(P_{e,1}^I, \dots, P_{e,n-1}^I, P_{e,n+1}^I, \dots, P_{e,N}^I) = \sum_{\substack{I \subseteq [N] \setminus \{n\} \\ |I|=L-1}} \prod_{i \in I} P_{e,i}^I \prod_{j \in [N] \setminus (I \cup \{n\})} (1 - P_{e,j}^I), \quad (2.12)$$

$$g_{L,N}^{II}(P_{e,1}^{II}, \dots, P_{e,n-1}^{II}, P_{e,n+1}^{II}, \dots, P_{e,N}^{II}) = \sum_{\substack{I \subseteq [N] \setminus \{n\} \\ |I|=N-L}} \prod_{i \in I} P_{e,i}^{II} \prod_{j \in [N] \setminus (I \cup \{n\})} (1 - P_{e,j}^{II}). \quad (2.13)$$

Then the optimal local decision rule (2.9) is rewritten to

$$\frac{f_{Y_n|H}(y_n|1)}{f_{Y_n|H}(y_n|0)} \underset{\hat{H}_n(y_n)=0}{\overset{\hat{H}_n(y_n)=1}{>}} \frac{(c_{10} - c_{00}) p_0 g_{L,N}^I(P_{e,1}^I, \dots, P_{e,n-1}^I, P_{e,n+1}^I, \dots, P_{e,N}^I)}{(c_{01} - c_{11}) (1 - p_0) g_{L,N}^{II}(P_{e,1}^{II}, \dots, P_{e,n-1}^{II}, P_{e,n+1}^{II}, \dots, P_{e,N}^{II})}. \quad (2.14)$$

While the optimal decision rule in the single-agent case is determined by the prior probabilities and the likelihood functions of observations, that in the distributed detection model with N agents is determined by the fusion rule as well. For fixed N , decrease of L leads to increase of the optimal decision threshold. On the other hand, for fixed L , decrease of N leads to decrease of the optimal decision threshold. If fewer decisions for 1 are required to make the global decision 1, agents are less likely to decide 1.

An extreme example can be found in psychological phenomena, the bystander effect. It refers to phenomenon that people do not help a victim, such as calling the police or an ambulance, when other people are nearby. The likelihood of help is inversely related to the number of bystanders. From the perspective of signal processing, calling the police is like the OR (1-out-of- N) fusion rule in the sense that just one call is sufficient to bring police officers to the scene. As N gets larger, people think that somebody will call the police even if they do not. The threshold in terms of emergency level is increased and the people get unlikely to call the police.

■ 2.3 Social Learning

Agents are not allowed to communicate with each other in the previous models. They can only send their local decisions to the fusion center, which takes responsibility of turning the local decisions into a global decision. However, if the agents can observe decisions made by other agents, then these observations will also give information about the right choice.

Consider multiple decision makers who do not perform distributed detection. Instead, they are detecting a hypothesis individually: Each agent takes the cost of her own decision and minimizes her own Bayes risk. If she can observe the choices

of some earlier-acting agents before making a decision, then she will try to learn from them and make a better decision. This learning behavior is called *social learning* or *observational learning*. The decisions observed by later-acting agents are called *public signals* and distinguished from *private signals*, which are observed only by individual agents like Y_n in Section 2.2.

The framework of sequential decision making with social learning was independently introduced in [17] and [18]. These works consider a sequence of agents who make decisions in a predetermined order. The agents also know the whole decision-making order and observe all decisions of previous agents. The papers observed that herding—all later-acting agents follow and never deviate from the public signals regardless of their private signals—can occur even if the public signals are incorrect.

Smith and Sørensen [19] showed that incorrect herding occurs with nonzero probability if private signals are boundedly informative.² However, agents will asymptotically settle on the optimal choice otherwise.

■ 2.3.1 Social Learning and Belief Update

There are N agents indexed as $1, 2, \dots, N$ corresponding to their decision-making orders. Agents observe iid private signals conditioned on H and all decisions of earlier-acting agents to perform binary hypothesis testing.

Consider decision making of Agent n . When she observes Y_n and $\widehat{H}_1, \widehat{H}_2, \dots, \widehat{H}_{n-1}$, her Bayes risk is

$$\begin{aligned} R_n &= \sum_{i=0}^1 \sum_{j=0}^1 c_{ji} \mathbb{P}\{\widehat{H} = j, H = i \mid \widehat{H}_1 = \widehat{h}_1, \widehat{H}_2 = \widehat{h}_2, \dots, \widehat{H}_{n-1} = \widehat{h}_{n-1}\} \\ &= \sum_{i=0}^1 \sum_{j=0}^1 c_{ji} p_{\widehat{H}_n | H}(j | i) p_{H | \widehat{H}_1, \dots, \widehat{H}_{n-1}}(i | \widehat{h}_1, \dots, \widehat{h}_{n-1}), \end{aligned}$$

where $p_{\widehat{H}_n | H, \widehat{H}_1, \dots, \widehat{H}_{n-1}}(\cdot | \cdot) = p_{\widehat{H}_n | H}(\cdot | \cdot)$ is used. This is true because the private signals are independent conditioned on H and so are the public signals. Hence, social learning

²A signal Y generated under H is called *boundedly informative* if there exists $\kappa > 0$ such that $\kappa < f_{Y|H}(y|h) < 1/\kappa$ for all y and h .

is equivalent to updating belief from $p_H(\cdot)$ to

$$p_{H|\widehat{H}_1, \dots, \widehat{H}_{n-1}}(\cdot | \widehat{h}_1, \dots, \widehat{h}_{n-1})$$

and performing the following likelihood ratio test:

$$\frac{f_{Y_n|H}(y_n|1)}{f_{Y_n|H}(y_n|0)} \underset{\widehat{H}_n(y_n)=0}{\overset{\widehat{H}_n(y_n)=1}{\gtrless}} \frac{(c_{10} - c_{00})p_{H|\widehat{H}_1, \dots, \widehat{H}_{n-1}}(0 | \widehat{h}_1, \dots, \widehat{h}_{n-1})}{(c_{01} - c_{11})p_{H|\widehat{H}_1, \dots, \widehat{H}_{n-1}}(1 | \widehat{h}_1, \dots, \widehat{h}_{n-1})}. \quad (2.15)$$

The likelihood ratio test (2.15) can also be obtained from (2.1) by treating all private and public signals as noisy observations:

$$\frac{f_{Y_n, \widehat{H}_1, \dots, \widehat{H}_{n-1}|H}(y_n, \widehat{h}_1, \dots, \widehat{h}_{n-1} | 1)}{f_{Y_n, \widehat{H}_1, \dots, \widehat{H}_{n-1}|H}(y_n, \widehat{h}_1, \dots, \widehat{h}_{n-1} | 0)} \underset{\widehat{H}_n(y_n)=0}{\overset{\widehat{H}_n(y_n)=1}{\gtrless}} \frac{(c_{10} - c_{00})p_0}{(c_{01} - c_{11})(1 - p_0)}.$$

■ 2.3.2 Herding Behavior

Suppose that the private signals are symmetric and binary-valued: For $0 < \epsilon < 0.5$,

$$\mathbb{P}\{Y_n = H\} = 1 - \epsilon \quad \text{and} \quad \mathbb{P}\{Y_n \neq H\} = \epsilon,$$

, $H \in \{0, 1\}$, $Y_n \in \{0, 1\}$ and $n = 1, 2, \dots, N$. The private signal Y_n indicates the true hypothesis H with high probability but also may indicate the wrong one with low probability.

The first agent performs the likelihood ratio test only with her private signal Y_1 ,

$$\frac{p_{Y_1|H}(y_1|1)}{p_{Y_1|H}(y_1|0)} \underset{\widehat{H}_1=0}{\overset{\widehat{H}_1=1}{\gtrless}} \frac{c_{10}p_0}{c_{01}(1 - p_0)}.$$

For meaningful discussion, it should be avoided that her decision is made regardless of her private signal. Thus, ϵ is assumed to satisfy

$$\frac{\epsilon}{1 - \epsilon} < \frac{c_{10}p_0}{c_{01}(1 - p_0)} < \frac{1 - \epsilon}{\epsilon}, \quad (2.16)$$

which is an identical condition to

$$\epsilon < \min \left\{ \frac{c_{10}p_0}{c_{10}p_0 + c_{01}(1-p_0)}, \frac{c_{01}(1-p_0)}{c_{10}p_0 + c_{01}(1-p_0)} \right\}.$$

Agent n updates her belief from p_0 to μ_n based on the public signals $\widehat{H}_1, \widehat{H}_2, \dots, \widehat{H}_{n-1}$:

$$\begin{aligned} \mu_n &= \mathbb{P}\{H = 0 \mid \widehat{H}_1 = \widehat{h}_1, \widehat{H}_2 = \widehat{h}_2, \dots, \widehat{H}_{n-1} = \widehat{h}_{n-1}\} \\ &= \frac{\mathbb{P}\{\widehat{H}_{n-1} = \widehat{h}_{n-1} \mid H = 0\} \mathbb{P}\{H = 0 \mid \widehat{H}_1 = \widehat{h}_1, \dots, \widehat{H}_{n-2} = \widehat{h}_{n-2}\}}{\sum_{h=0}^1 \mathbb{P}\{\widehat{H}_{n-1} = \widehat{h}_{n-1} \mid H = h\} \mathbb{P}\{H = h \mid \widehat{H}_1 = \widehat{h}_1, \dots, \widehat{H}_{n-2} = \widehat{h}_{n-2}\}} \\ &= \frac{\mathbb{P}\{\widehat{H}_{n-1} = \widehat{h}_{n-1} \mid H = 0\} \mu_{n-1}}{\mathbb{P}\{\widehat{H}_{n-1} = \widehat{h}_{n-1} \mid H = 0\} \mu_{n-1} + \mathbb{P}\{\widehat{H}_{n-1} = \widehat{h}_{n-1} \mid H = 1\} (1 - \mu_{n-1})}, \end{aligned} \quad (2.17)$$

with the boundary condition $\mu_1 = p_0$. The equation (2.17) implies that the updated belief is public information because all agents who observe the public signals would update their belief identically. In addition, $\mu_n \geq \mu_{n-1}$ if $\widehat{H}_{n-1} = 0$, and $\mu_n \leq \mu_{n-1}$ if $\widehat{H}_{n-1} = 1$.

Without loss of generality, let us assume that the first agent chooses 0. The second agent observes it and can infer that the private signal of the first agent is $Y_1 = 0$. The belief is updated to

$$\mu_2 = \frac{p_0(1-\epsilon)}{p_0(1-\epsilon) + (1-p_0)\epsilon}.$$

The assumption (2.16) leads to the following inequality

$$\frac{c_{10}\mu_2}{c_{01}(1-\mu_2)} > \frac{c_{10}p_0}{c_{01}(1-p_0)} > \frac{\epsilon}{1-\epsilon},$$

and there are two possibilities for μ_2 :

- When $\frac{c_{10}\mu_2}{c_{01}(1-\mu_2)} \geq \frac{1-\epsilon}{\epsilon}$: The second agent will always choose 0 regardless of her private signal. Even if her private signal is $Y_2 = 1$, the likelihood ratio of Y_2 is

$$\frac{p_{Y_2|H}(1|1)}{p_{Y_2|H}(1|0)} = \frac{1-\epsilon}{\epsilon} < \frac{c_{10}\mu_2}{c_{01}(1-\mu_2)}.$$

Furthermore, all later-acting agents will always choose 0 because μ_n will not be smaller than μ_2 for $n > 2$.

- When $\frac{\epsilon}{1-\epsilon} < \frac{c_{10}\mu_2}{c_{01}(1-\mu_2)} < \frac{1-\epsilon}{\epsilon}$: The second agent will choose according to her private signal Y_2 . If she has $Y_2 = 1$ then she will choose 1 and μ_3 will be smaller than μ_2 . Thus, the third agent will also choose according to her private signal.

If the second agent observes $Y_2 = 0$, however, then she chooses 0. The third agent can infer from the public signals $\widehat{H}_1 = \widehat{H}_2 = 0$ that the first two agents have observed $Y_1 = 0$ and $Y_2 = 0$, respectively. The belief is updated to

$$\mu_3 = \frac{p_0(1-\epsilon)^2}{p_0(1-\epsilon)^2 + (1-p_0)\epsilon^2}.$$

When she observes $Y_3 = 1$,

$$\begin{aligned} \frac{p_{Y_3|H}(1|1)}{p_{Y_3|H}(1|0)} &= \frac{1-\epsilon}{\epsilon} \\ &< \frac{1-\epsilon}{\epsilon} \times \left(\frac{c_{10}p_0}{c_{01}(1-p_0)} \frac{1-\epsilon}{\epsilon} \right) = \frac{c_{10}p_0(1-\epsilon)^2}{c_{01}(1-p_0)\epsilon^2} = \frac{c_{10}\mu_3}{c_{01}(1-\mu_3)}, \end{aligned}$$

where the inequality comes from (2.16). In other words, the belief becomes too strong for the third agent to reject the public signals. Thus, she will adopt the decisions of the first two agents and so will later-acting agents.

In conclusion, the belief is updated based on each public signal. Once it goes above a boundary, herding occurs. In the binary-symmetric-private-signal case, the boundary is given as follows:

- Agents herd on 0 if $\mu_n \geq \frac{p_0(1-\epsilon)^2}{p_0(1-\epsilon)^2 + (1-p_0)\epsilon^2}$.
- Agents herd on 1 if $\mu_n \leq \frac{p_0\epsilon^2}{p_0\epsilon^2 + (1-p_0)(1-\epsilon)^2}$ from the symmetry.

The herding occurs when at least the first two agents coincide with their decisions. If both of them are wrong, whose probability is $\epsilon^2 > 0$, all the later-acting agents happen to adopt wrong decisions. Incorrect herding occurs with probability higher than ϵ^2 because it also occurs, for example, when $\widehat{H}_1 = 0, \widehat{H}_2 = 1, \widehat{H}_3 = 0$, and $\widehat{H}_4 = 0$.

When the private signal is boundedly informative, herding occurs when the belief becomes too strong. On the other hand, when the private signal is unbounded, it

can be strong enough to reject the public signals with positive probability. The later-acting agents can asymptotically converge to the true hypothesis [19]. This thesis does not restrict the private signals to be boundedly informative.

Human Perception of Prior Probability

Economic theory understands human behavior with the fundamental assumption that humans act rationally through the optimization of their expected payoff. The payoff is expressed in terms of the amounts of money gained or lost in many cases but can include terms of other quantities, such as time [20]. Furthermore, the payoff does not have to be quantifiable. It includes regret [21, 22], disappointment [23], ambiguity aversion [24], and fairness [25].

However, a human's cognitive process is different from the calculation process of a machine. The optimization of the payoff should involve not only the computational power but also the ability to anticipate and quantify exact outcomes and their likelihoods. The complex decision-making process is hardly observed in human choice situations, and Simon [26] suggested modification of procedures: simple payoff functions, costly information gathering, and lack of a complete order of priority when the outcome of the payoff function is a vector. With observing that human decision making is to satisfice rather to optimize, he thought that "a great deal can be learned about rational decision making by taking into account, at the outset, the limitations upon the capacities and complexity of the organism, and by taking account of the fact that the environments to which it must adapt possess properties that permit further simplification of its choice mechanisms" [27].

The notion of *bounded rationality* contains this idea. The models of bounded rationality describe more than the outcome of a decision. They also describe how hu-

mans heuristically develop approximate mechanisms and how human decision-making processes reach a judgment or decision.

This thesis considers human cognitive limitations but still allows optimization. It follows Marschak's approach, which is termed *costly rationality* [7, 28]. Under costly rationality, a rational agent optimizes the payoff with taking into account costs of activities of decision making, such as observation, computation, memory, and communication.

The constraint that is taken into account in the thesis is imperfect perception of prior probability. Section 3.1 reviews psychological literature about the effect of attractiveness on human decision making and the perception of prior probability. Section 3.2 describes a mathematical abstraction of the perception process and estimates the prior probability.

■ 3.1 Studies on Human Decision Making

Rational human decision making is known to resemble Bayesian reasoning [1–4]. However, not all human decision making seems rational. One of the irrational behaviors is the human tendency to discriminate against unattractive people. Even though a person's appearance is unrelated to the person's emotions, intelligence, or behaviors, many psychological studies have reported evidence of biases based on physical attractiveness.

Stewart [29] studied the effect of attractiveness of a defendant on the severity of punishment. He had 10 observers who visited 74 criminal trials and rated the defendants' attractiveness on 7-point bipolar scale. The observers attempted to visit whenever possible; they watched the defendant for 30 minutes and rated the attractiveness on given standard rating forms. The forms have 9 items, such as richness and education as well as attractiveness, so that the observers could not know that the experiment would only care about attractiveness of the defendants.

He analyzed all data observed from 73 trials and found the relation of the attractiveness rates, race, seriousness of the crime, and conviction/acquittal. The minimum

and maximum sentences were strongly inversely related to the attractiveness. The seriousness of crime¹ was also negatively correlated to the attractiveness. When the seriousness was controlled in order to figure out the causality of severity of sentence, the correlation between the attractiveness and the severity of sentence was reduced but existed.

Similarly, Efran [30] tested a hypothesis that physically attractive defendants would be more positively evaluated than unattractive ones. The results of a simulated jury task supported that attractive defendants were evaluated with less probability of guilty verdict and less severe recommended punishment than unattractive defendants were even though appearance should be irrelevant to judicial decisions. Also, similar results were observed in [31].

These results may imply that jurors just make a more favorable verdict to a more attractive defendant. Stephan and Tully [32] studied attractiveness of plaintiffs and designed a survey with eight personal injury suits each brought by an individual against another as a result of an automobile accident. The result supported their hypothesis that a physically attractive plaintiff is favored over an unattractive plaintiff in assessing liability or in the amount of money awarded. Considering that the attractiveness of plaintiff is truly irrelevant to assessing damage from the accident, these studies seem to indicate that “what is beautiful is good.”

In fact, the advantages of attractiveness were studied from many other aspects as well. Dion et al. revealed that a physical attractiveness stereotype exists [33]. Their experiment showed that attractive men and women were expected to attain more reputable occupations than less-attractive people were. Furthermore, attractive people were expected to experience happier marriages, be better spouses and parents, and have happier social and professional lives.

Furthermore, a child’s attractiveness is also significantly associated with a teacher’s expectations about the child’s intelligence, parents’ interest, educational potential,

¹The crimes were categorized into three groups with respect to their seriousness: The most serious crimes include murder, voluntary manslaughter, and rape. The second group consists of armed robbery, robbery, burglary, aggravated assault, and involuntary manslaughter. The least serious crimes were theft by taking, deception, victimless crimes, and minor drug offenses.

and social potential [34]. A physically attractive person is expected to be more socially skillful and likable [35]. An attractive person is evaluated to be more suitable for employment and more likely to advance [36]. Even experienced managers have the attractiveness biases even though the biases tend to decrease as managerial experience increases.

The biased punishment based on physical attractiveness may also result from the jurors' recognition that a more attractive defendant seems less likely to commit a crime. Sigall and Ostrove [37] looked for a cognitive explanation of the biases. They studied whether the physical attractiveness of a criminal defendant has a different effect on a crime unrelated to attractiveness (i.e., burglary) and on an attractiveness-related one (i.e., swindle). Their experiment revealed that an unattractive defendant was more severely punished when the crime was attractiveness-unrelated, but an attractive defendant was treated more harshly when the crime was attractiveness-related. However, since the difference in sentences assigned to an attractive and an unattractive defendant was not statistically significant, the authors conservatively concluded that the advantages of attractiveness are lost when the crime is attractiveness-related.

Smith and Hed [38] reviewed the correlation between attractiveness and the severity of sentence studied in [37]. They formed a group of three jurors to figure out the difference between a group and an individual decision maker—whether the jurors unfavorably judge an attractive defendant in a swindle case even after a discussion as individual jurors do without a discussion. They found that, even though in burglary cases attractive defendants were sentenced significantly less harshly than unattractive defendants were, in swindle cases attractive defendants were sentenced slightly but not significantly more harshly than unattractive defendants were. These results coincide with the results in [37] in which judgments were made by individuals.

Darby and Jeffers [5] also observed that more attractive defendants are less frequently convicted and less severely punished. Besides, they asked their experiment subjects to evaluate defendants' likability, trustworthiness, and responsibility for charges. Attractive defendants were rated as more likable, more trustworthy, and

less responsible for the offense. The results imply that the biased verdict may be related to jurors' perception that the attractive defendant is less likely to commit the charged crime.

The relation is more obvious in [6]. Brown et al. studied the effect of eyeglasses on juror decisions in violent crimes. They not only looked at the verdicts with and without eyeglasses but also asked subjects to rate defendants' intelligence, attractiveness, friendliness, and physical threateningness on a scale of 1 to 7.

People turned out to feel a defendant more intelligent and less threatening statistically significantly when the defendant was wearing eyeglasses. It was found that this effect of eyeglasses led to fewer guilty verdict for the defendant with eyeglasses than for one without eyeglasses.

To summarize, the psychological studies have verified that the attractiveness stereotype does exist: People are favorable to a physically attractive person. However, it is not the only reason that attractive defendants are less severely punished than unattractive defendants are. It is also because people think that attractive people would be less likely to commit a crime. Otherwise, attractive defendants would have been also less harshly punished in the cases of attractiveness-related crimes. This observation is a motivation of the thesis.

■ 3.2 Estimation of Prior Probabilities

Again, it turned out that defendants wearing eyeglasses received fewer guilty verdicts than ones without eyeglasses did [6]. From a logical standpoint, there should not be such a correlation. Eyeglasses can neither help nor prevent committing a crime. Furthermore, jurors must be doing their best to make fair and reasonable decisions based on evidence.

The correlation can be found from Bayesian reasoning, which requires prior probability to be performed. If a trial is likened to a hypothesis testing problem, the hypotheses are whether a defendant is guilty or innocent. Observed signals are the evidence presented by a prosecutor or defense counsel. Jurors are decision-making




Figure 3-1. An extensive hypothesis testing model proposed in the thesis.

agents that make a verdict upon observing the signals. One element of hypothesis testing missing here is the prior probability that the defendant commits the charged crime.

The prior probability matters to the verdict and is critical especially when the evidence is ambiguous. The problem is that the jurors cannot know the defendant's prior probability. Hence, they judge the probability based on the defendant's appearance. They evaluate the defendant's personality, such as intelligence, attractiveness, friendliness, and dangerousness, and infer the prior probability.

In order to capture practical decision-making agents like human beings, this thesis proposes a modification: treat the prior probability p_0 as a hyperparameter with a distribution $f_{P_0}(p_0)$, Figure 3-1. A state H , the subject of detection, is drawn from a Bernoulli distribution with parameter p_0 but the agent does not know the value of p_0 , unlike in general hypothesis testing models. The prior probability p_0 as well as the state H are hidden (unknown) variables. Instead of p_0 , the agent can observe a signal $T = t$, which is correlated with p_0 by $f_{T|P_0}(t|p_0)$. The agent performs hypothesis testing with two observations $T = t$ and $Y = y$. Note that the problem becomes the same as the classical binary hypothesis testing if T is always equal to p_0 .

The objective of this problem is not to estimate p_0 accurately but to detect H accurately. The criterion is the following Bayes risk:

$$\begin{aligned}
 R &= c_{10}\mathbb{P}\{\widehat{H} = 1, H = 0\} + c_{01}\mathbb{P}\{\widehat{H} = 0, H = 1\} \\
 &= c_{10} \int_{\mathcal{Y}_1} f_{T,Y|H}(t, y|0)P_H(0) dt dy + c_{01} \int_{\mathcal{Y}_0} f_{T,Y|H}(t, y|1)P_H(1) dt dy,
 \end{aligned}$$

where \mathcal{Y}_0 and \mathcal{Y}_1 here are defined over two-dimensional space of Y and T . The regions

\mathcal{Y}_0 and \mathcal{Y}_1 that minimize the Bayes risk are given by

$$\begin{cases} (t, y) \in \mathcal{Y}_1, & \text{if } c_{10}f_{T,Y|H}(t, y|0)P_H(0) \leq c_{01}f_{T,Y|H}(t, y|1)P_H(1), \\ (t, y) \in \mathcal{Y}_0, & \text{otherwise.} \end{cases}$$

The corresponding optimal decision rule $\widehat{H}(t, y)$ is the likelihood ratio test with two observations $T = t$ and $Y = y$,

$$\frac{f_{T,Y|H}(t, y|1)}{f_{T,Y|H}(t, y|0)} \underset{\widehat{H}(t,y)=0}{\overset{\widehat{H}(t,y)=1}{\gtrless}} \frac{c_{10}p_H(0)}{c_{01}p_H(1)}. \quad (3.1)$$

The signals T and Y are correlated but they are independent conditioned on H ,

$$f_{T,Y|H}(t, y|h) = f_{T|H}(t|h)f_{Y|H}(y|h).$$

The conditional independence changes (3.1) to

$$\frac{f_{Y|H}(y|1)}{f_{Y|H}(y|0)} \underset{\widehat{h}(t,y)=0}{\overset{\widehat{h}(t,y)=1}{\gtrless}} \frac{c_{10}p_H(0)f_{T|H}(y|0)}{c_{01}p_H(1)f_{T|H}(y|1)}. \quad (3.2)$$

At the right-hand side of (3.2),

$$\begin{aligned} p_H(0)f_{T|H}(t|0) &= f_{H,T}(0, T) = \int_0^1 f_{H,T,P_0}(0, t, p_0) dp_0 \\ &\stackrel{(a)}{=} \int_0^1 p_{H|P_0}(0|p_0)f_{T,p_0}(t, p_0) dp_0 \\ &\stackrel{(b)}{=} \int_0^1 p_0 f_{P_0|T}(p_0|t)f_T(t) dp \\ &= \mathbb{E}[P_0|T = t]f_T(t), \end{aligned} \quad (3.3)$$

where the equality (a) holds because H and T are independent conditioned on P_0 ; the equality (b) holds because $p_{H|P_0}(0|p_0) = p_0$ by definition. Likewise, $p_H(1)f_{T|H}(t|1) = (1 - \mathbb{E}[P|T = t])f_T(t)$. The decision rule (3.2) becomes

$$\frac{f_{Y|H}(y|1)}{f_{Y|H}(y|0)} \underset{\widehat{h}(t,y)=0}{\overset{\widehat{h}(t,y)=1}{\gtrless}} = \frac{c_{10}\mathbb{E}[P_0|T = t]}{c_{01}(1 - \mathbb{E}[P_0|T = t])}, \quad (3.4)$$

which looks similar to (2.1). It turns out that the optimal decision rule is the likelihood ratio test with the observation $Y = y$ when the agent estimates the prior probability at $\mathbb{E}[P_0|T = t]$.

Therefore, the decision-making process can be split into two steps: estimating the prior probability upon observing T and detecting the true state upon observing Y . This result is useful because T and Y can be dealt with separately. It restores our problem to the classical hypothesis testing problem except that the prior probability used by the Bayesian agent is the conditional mean of prior probability $\mathbb{E}[P_0|T = t]$.

In reality, people would know the likelihood $f_{P_0|T}$ differently. Furthermore, they may use other kinds of estimators. Therefore human decision makers are more likely to differently perceive the prior probability even if they observe the same signal T . In the jury example [6], every juror can watch the defendant wearing eyeglasses. However, each juror would have a different feeling of how violent the defendant looks and different guesses of the defendant's prior probability of guilt.

It is too complex to assume the likelihood functions $f_{P_0|T}$ differently for each agent and consider their estimation processes. Instead of the whole process of estimating prior probability, let us assume that agents have already estimated the prior probability and are ready to perform hypothesis testing. Through most of the thesis, agents are assumed to make decisions with individually perceived prior probabilities (or perceived *belief*).

Distributed Detection with Symmetric Fusion

One of the simplest ways to decrease the Bayes risk is to form a decision-making team. A team considered in the thesis consists of N decision-making agents who individually observe conditionally independent signals. Even though it leads to the best result to directly integrate the signals, it causes high communication costs.

The advantage of the multiplicity is still effective even when agents locally make decisions, i.e., quantize their observations to binary values, then fuse the decisions. Suppose that the private signals are corrupted by iid additive Gaussian noises. When all agents are correctly aware of the prior probability, the team decision making with a symmetric L -out-of- N fusion rule is equivalent to a single decision making with the L th largest signal $Y_{(L)}$ [12]. The pdf of the L th largest noise is computed as follows according to the order statistics [39]:

$$f_{W_{(L)}}(w) = \frac{N!}{(N-L)!(L-1)!} [F_W(w)]^{(N-L)} [1 - F_W(w)]^{(L-1)} f_W(w). \quad (4.1)$$

For odd number N and the majority fusion rule ($L = (N + 1)/2$), for instance, the effective variance of $W_{(L)}$ is listed in Table 4.1 [40].

The advantage described in Table 4.1 is fully achieved only by the perfect agents who know the prior probability. The imperfect agents under our assumption, who do not know the prior probability, would be outperformed by them. This chapter discusses the limited performance of teams of imperfect agents.

N	L	Variance	N	L	Variance
1	1	1	11	6	0.1372
3	2	0.4487	13	7	0.1168
5	3	0.2868	15	8	0.1017
7	4	0.2104	17	9	0.0900
9	5	0.1661	19	10	0.0808

Table 4.1. Variance of the median of N iid Gaussian random variables




Figure 4-1. An agent observes N conditionally independent signals Y_1, Y_2, \dots, Y_N .

Section 4.1 evaluates an imperfect agent. Section 4.2 investigates the performance of a team of imperfect agents focusing on the comparison between teams of identical agents and diverse ones. Section 4.3 shows that imperfect agents can outperform perfect agents if they do observe non-identically distributed signals. Section 4.4 concludes the chapter.

■ 4.1 An Imperfect Agent

Consider a hypothesis testing problem of a single agent. The agent detects a state H , whose prior probability is $p_0 = \mathbb{P}\{H = 0\}$ and $p_1 = \mathbb{P}\{H = 1\} = 1 - p_0$. However she is imperfect and does not know p_0 . Instead, she perceives the prior probability $\mathbb{P}\{H = 0\}$ as $q \neq p_0$. Let's compare her performance to a perfect agent's.

■ 4.1.1 The Number of Observations

Suppose that the imperfect agent can observe N conditionally independent signals Y_1, \dots, Y_N , Figure 4-1. The likelihoods of her observations are expressed as the mul-

tiplication of the likelihoods of individual observations:

$$P_{Y_1, \dots, Y_N | H}(y_1, \dots, y_N | 0) = \prod_{n=1}^N P_{Y_n | H}(y_n | 0),$$

$$P_{Y_1, \dots, Y_N | H}(y_1, \dots, y_N | 1) = \prod_{n=1}^N P_{Y_n | H}(y_n | 1).$$

With these likelihoods and her perceived prior belief, she would perform the following LRT:

$$\frac{\prod_{n=1}^N P_{Y_n | H}(y_n | 1)}{\prod_{n=1}^N P_{Y_n | H}(y_n | 0)} \underset{\widehat{H}_q(y_1, \dots, y_N)=1}{\underset{\widehat{H}_q(y_1, \dots, y_N)=0}{\geq}} \frac{c_{10}q}{c_{01}(1-q)}, \quad (4.2)$$

where $\widehat{H}_q(\cdot)$ denotes the decision rule optimized for prior probability q . This LRT determines her false alarm probability $\mathbb{P}\{\widehat{H}_q(y_1, \dots, y_N) = 1 | H = 0\}$ and missed detection probability $\mathbb{P}\{\widehat{H}_q(y_1, \dots, y_N) = 0 | H = 1\}$.

If the observed signals are $Y_n = H + W_n$, where W_n are iid Gaussian random variables with zero mean and unit variance, then observing the N signals is equivalent to observing a signal corrupted by additive Gaussian noise with zero mean and variance $1/N$. Let N_{ob} denote the minimum number of observations for her to perform no worse than a perfect agent observing one signal:

$$N_{\text{ob}} = \min N$$

$$\text{s.t. } c_{10}p_0\mathbb{P}\{\widehat{H}_q(y_1, \dots, y_N) = 0 | H = 1\} + c_{10}p_1\mathbb{P}\{\widehat{H}_q(y_1, \dots, y_N) = 1 | H = 0\}$$

$$\leq c_{10}p_0\mathbb{P}\{\widehat{H}_{p_0}(y_1) = 0 | H = 1\} + c_{01}p_1\mathbb{P}\{\widehat{H}_{p_0}(y_1) = 1 | H = 0\}. \quad (4.3)$$

The value of N_{ob} depends on the true prior probability p_0 , the perceived prior belief q , and the Bayes costs c_{10} and c_{01} . For instance, when the prior probability of a decision-making subject is $p_0 = p_1 = 0.5$, N_{ob} for perceived belief $q \in [0.01, 0.99]$ is depicted in Figure 4-2. The worst perceived belief ($q = 0$ or $q = 1$) cannot be overcome even by infinite observations. However, if she reasonably perceives the prior probability, e.g., within an error of 0.3, then only 3 or less observations are sufficient. Figure 4-2 implies that having a wrong belief 0.25 when the true prior




Figure 4-2. The minimum number of observations for an agent, who perceives prior probability as $q \in [0.01, 0.99]$, to reduce the expected cost below that of a perfect agent who knows p_0 . All observations are $Y_i = H + W_i$, where $W_i \sim \text{iid } \mathcal{N}(0, 1)$. $p_0 = 0.5$ and $c_{10} = c_{01} = 1$.

probability is 0.5 is similar to observing a signal with half of the power of the perfect agent's signal.

■ 4.1.2 The Number of Agents

Now suppose that there are N imperfect agents who perceive the prior probability as q , Figure 4-3. Agent n observes a conditionally iid signal $Y_n = y_n$ and makes a local decision. The local decisions are fused by a known L -out-of- N fusion rule.

When the agents use decision thresholds $\lambda_1, \lambda_2, \dots, \lambda_N$, respectively, their global Bayes risk is

$$\begin{aligned}
 R &= c_{10}p_0P_E^I + c_{01}p_1P_E^{II} \\
 &= c_{10}p_0G_{L,N}^I(P_e^I(\lambda_1), \dots, P_e^I(\lambda_N)) + c_{01}p_1G_{L,N}^{II}(P_e^{II}(\lambda_1), \dots, P_e^{II}(\lambda_N)),
 \end{aligned}$$

where the functions $G_{L,N}^I$ and $G_{L,N}^{II}$ are defined in (2.10) and (2.11). Let us restrict the agents to use the same decision thresholds, i.e., $\lambda_1 = \lambda_2 = \dots = \lambda_N = \lambda$ because the restriction causes little or no loss of performance [16]. Then their decision threshold




Figure 4-3. Multiple agents individually perform hypothesis testing then a fusion center fuses their decisions.

λ is the solution to (2.14):

$$\frac{P_{Y_n|H}(\lambda|1)}{P_{Y_n|H}(\lambda|0)} = \frac{c_{10}qg_{L,N}^I(P_e^I(\lambda), \dots, P_e^I(\lambda))}{c_{01}(1-q)g_{L,N}^{II}(P_e^{II}(\lambda), \dots, P_e^{II}(\lambda))}.$$

Each agent performs the following comparison to make her local decision:

$$y_n \underset{\hat{H}_q(y_n)=0}{\overset{\hat{H}_q(y_n)=1}{\geq}} \lambda.$$

Individual imperfect agents are outperformed by the perfect agent but, as a team, the imperfect agents can outperform her. Let N_{ag} denote the minimum number of imperfect agents who can together outperform the perfect agent:

$$\begin{aligned} N_{\text{ag}} &= \min N \\ \text{s.t. } \min_L &\{c_{10}p_0G_{L,N}^I(P_e^I(\lambda), \dots, P_e^I(\lambda)) + c_{01}p_1G_{L,N}^{II}(P_e^{II}(\lambda), \dots, P_e^{II}(\lambda))\} \\ &\leq c_{10}p_0\mathbb{P}\{\hat{H}_{p_0}(y_1) = 0 | H = 1\} + c_{01}p_1\mathbb{P}\{\hat{H}_{p_0}(y_1) = 1 | H = 0\}. \end{aligned}$$

Figure 4-4 depicts the minimum number of imperfect agents who can outperform a single perfect agent. Compared to Figure 4-2, Figure 4-4 shows that the number of required agents is larger than the number of observations required. It is because the




Figure 4-4. The minimum number of imperfect agents, who perceive prior probability as $q \in [0.01, 0.99]$, to reduce the expected cost below that of a perfect agent who knows p_0 . All agents observe $Y_i = H + W_i$, where $W_i \sim \text{iid } \mathcal{N}(0, 1)$. $p_0 = 0.5$ and $c_{10} = c_{01} = 1$.

agents lose information when they make local decisions and the L -out-of- N rule may not be the optimal fusion rule.

■ 4.2 Team of Imperfect Agents

As discussed in Section 4.1, one way to improve decision making is to observe N signals; another way is to group N agents so that they make a decision as a team. The latter case is identical to the former if all agents send their observations to their fusion center. However, it requires communication channels of infinite capacity between agents and the fusion center.

Thus, the agents are constrained in the thesis to be able to send a 1-bit signal to the fusion center. With exchange of low communication costs, performance will be degraded due to loss of information and ignorance of prior probability. This section discusses their team performance focusing on the comparison between teams of identical agents and diverse ones.

■ 4.2.1 A Team of Agents Who Observe Conditionally IID Signals

First, let us compare a team of N diverse agents and a team of N identical agents. A team of identical agents is already discussed in Section 4.1.2. Diverse agents are defined as agents who differently perceive the prior probability as $q_1, q_2 \dots, q_N$, respectively. The model of diverse team is the same as that of an identical team except that the agents have different beliefs.

To recapitulate decision making by identical agents, their decision threshold is the solution to the equation

$$\frac{P_{Y_n|H}(\lambda|1)}{P_{Y_n|H}(\lambda|0)} = \frac{c_{10}qg_{L,N}^I(P_e^I(\lambda), \dots, P_e^I(\lambda))}{c_{01}(1-q)g_{L,N}^{II}(P_e^{II}(\lambda), \dots, P_e^{II}(\lambda))}. \quad (4.4)$$

Let $r_{L,N} : [0, 1] \mapsto (-\infty, \infty)$ denote the mapping from the belief q to the decision threshold λ . In the process of determining the decision-making rule, it is implied that the agent and other colleague agents agree that the prior probability is q . All agents use the same equation (4.4) to determine their decision thresholds and their decision thresholds are identical.

The agents in the diverse team also know the number of decision makers and the fusion rule. Thus, they determine their decision thresholds from an equation of the same form as (4.4). However, contrary to the identical agents who know others' beliefs even if they do not communicate them (because they all have the same beliefs), these agents do not know others' beliefs if they do not communicate them.

Since Agent n does not know what beliefs other agents have, she just assumes that the other agents would have the same belief as hers. Accordingly, she assumes that the other agents also use the same decision threshold λ_n and will determine her decision threshold from the equation

$$\frac{P_{Y_n|H}(\lambda_n|1)}{P_{Y_n|H}(\lambda_n|0)} = \frac{c_{10}q_n g_{L,N}^I(P_e^I(\lambda_n), \dots, P_e^I(\lambda_n))}{c_{01}(1-q_n)g_{L,N}^{II}(P_e^{II}(\lambda_n), \dots, P_e^{II}(\lambda_n))}, \quad (4.5)$$

i.e., $\lambda_n = r_{L,N}(q_n)$. This rule only holds for Agent n . Agent i will use the threshold $\lambda_i = r_{L,N}(q_i)$, which is different from λ_n if $q_i \neq q_n$. Note that λ_n does not depend




Figure 4-5. The Bayes risks of a team of individual agents ($q_1 = q_2$) and a team of diverse agents ($q_1 \neq q_2$) for 1-out-of-2 fusion rule. $c_{10} = c_{01} = 1$. (a) $q = 0.4$, $q_1 = 0.4$, $q_2 = 0.6$. (b) $q = 0.5$, $q_1 = 0.5$, $q_2 = 0.6$.

on other agents' perceived beliefs. No matter whom she composes a team with, her decision rule is the same as long as her perceived belief, L , and N remain the same.

For the identical team, the performance is measured by the team Bayes risk,

$$\begin{aligned}
 R_I &= c_{10}p_0P_E^I + c_{01}p_1P_E^{II} \\
 &= c_{10}p_0G_{L,N}^I(P_e^I(\lambda), \dots, P_e^I(\lambda)) + c_{01}p_1G_{L,N}^{II}(P_e^{II}(\lambda), \dots, P_e^{II}(\lambda)). \quad (4.6)
 \end{aligned}$$

The performance of the diverse team is also measured by the same formula except that the agents have different decision-making rules:

$$\begin{aligned}
 R_D &= c_{10}p_0P_E^I + c_{01}p_1P_E^{II} \\
 &= c_{10}p_0G_{L,N}^I(P_e^I(\lambda_1), \dots, P_e^I(\lambda_N)) + c_{01}p_1G_{L,N}^{II}(P_e^{II}(\lambda_1), \dots, P_e^{II}(\lambda_N)). \quad (4.7)
 \end{aligned}$$

Comparison of the performances (4.6) and (4.7) are highly dependent on the agents beliefs q, q_1, q_2, \dots, q_N as well as the prior probability p_0 . For example, Figure 4-5a depicts the Bayes risks of a team of two identical agents who perceive the prior probability as 0.4 and a team of two diverse agents who respectively perceive it as 0.4 and 0.6. Their Bayes risks are linear in p_0 because the agents' decision thresholds are

constant for p_0 . The identical team performs better for small p_0 because Agent 2 in the diverse team has a belief with a larger error ($|0.6 - p_0| > |0.4 - p_0|$). On the contrary, the diverse team performs better for large p_0 because all agents in the identical team have beliefs with larger errors ($|0.4 - p_0| > |0.6 - p_0|$).

An interesting comparison can be made with regard to stability of the team. Suppose that the agents who perceived prior probability as 0.4 have a new belief 0.5 as in Figure 4-5b. The slopes of Bayes risk change because of the change of their beliefs. The slope of the identical team changes from 0.2406 to 0.03124 while that of the diverse team changes from 0.05633 to -0.07159. The difference is larger for the identical team than for the diverse team. Since the slope is equal to $c_{10}P_E^I - c_{01}P_E^{II}$, the identical team experiences more radical changes of error probabilities.

The following theorems give a rough comparison of the stability.

Theorem 4.1. *When a team of N agents who respectively have beliefs q_1, q_2, \dots, q_N make a decision with the L -out-of- N fusion rule, their global false alarm probability is given by $P_E^I = G_{L,N}^I(P_e^I(r_{L,N}(q_1)), P_e^I(r_{L,N}(q_2)), \dots, P_e^I(r_{L,N}(q_N)))$. Let us use a simple notation $P_E^I = E_{L,N}^I(q_1, q_2, \dots, q_N)$ for the global false alarm probability. Then for any $n \leq N$ and any beliefs q_0, q_1, \dots, q_{N-n} ,*

$$\left. \frac{d}{dq} E_{L,N}^I(\underbrace{q, \dots, q}_n, q_1, \dots, q_{N-n}) \right|_{q=q_0} = \frac{n}{n-1} \frac{d}{dq} E_{L,N}^I(\underbrace{q, \dots, q}_{n-1}, q_0, q_1, \dots, q_{N-n}) \Big|_{q=q_0} . \quad (4.8)$$

In other words, the change of false alarm probability is proportional to the number of agents who change their beliefs.

Proof. The proof is in Appendix 4.A. □

Theorem 4.2. *When a team of N agents who respectively have beliefs q_1, q_2, \dots, q_N make a decision with the L -out-of- N fusion rule, their global missed detection probability is given by $P_E^{II} = G_{L,N}^{II}(P_e^I(r_{L,N}(q_1)), P_e^I(r_{L,N}(q_2)), \dots, P_e^I(r_{L,N}(q_N)))$. Let us use a simple notation $P_E^{II} = E_{L,N}^{II}(q_1, q_2, \dots, q_N)$ for the global missed detection probability.*

Then for any $n \leq N$ and any beliefs q_0, q_1, \dots, q_{N-n} ,

$$\left. \frac{d}{dq} E_{L,N}^{\text{II}}(q, \dots, q, \underbrace{q_1, \dots, q_{N-n}}_n) \right|_{q=q_0} = \frac{n}{n-1} \frac{d}{dq} E_{L,N}^{\text{II}}(q, \dots, q, \underbrace{q_0, q_1, \dots, q_{N-n}}_{n-1}) \Big|_{q=q_0} . \quad (4.9)$$

Proof. Due to the symmetry, this theorem is proven by the proof of Theorem 4.1 in Appendix 4.A with change of superscripts I to II and L to $N - L + 1$. \square

Corollary 4.3. For any $m < n \leq N$ and any beliefs q_0, q_1, \dots, q_{N-m} ,

$$\begin{aligned} \left. \frac{d}{dq} E_{L,N}^{\text{I}}(q, \dots, q, \underbrace{q_1, \dots, q_{N-n}}_n) \right|_{q=q_0} &= \frac{n}{m} \frac{d}{dq} E_{L,N}^{\text{I}}(\underbrace{q, \dots, q}_m, \underbrace{q_0, \dots, q_0}_{n-m}, q_1, \dots, q_{N-n}) \Big|_{q=q_0} , \\ \left. \frac{d}{dq} E_{L,N}^{\text{II}}(q, \dots, q, \underbrace{q_1, \dots, q_{N-n}}_n) \right|_{q=q_0} &= \frac{n}{m} \frac{d}{dq} E_{L,N}^{\text{II}}(\underbrace{q, \dots, q}_m, \underbrace{q_0, \dots, q_0}_{n-m}, q_1, \dots, q_{N-n}) \Big|_{q=q_0} . \end{aligned}$$

Proof. For the false alarm probability,

$$\begin{aligned} &\left. \frac{d}{dq} E_{L,N}^{\text{I}}(q, \dots, q, \underbrace{q_1, \dots, q_{N-n}}_n) \right|_{q=q_0} \\ &= \frac{n}{n-1} \times \left. \frac{d}{dq} E_{L,N}^{\text{I}}(q, \dots, q, \underbrace{q_0, q_1, \dots, q_{N-n}}_{n-1}) \right|_{q=q_0} \\ &= \frac{n}{n-1} \times \frac{n-1}{n-2} \times \left. \frac{d}{dq} E_{L,N}^{\text{I}}(\underbrace{q, \dots, q}_{n-2}, \underbrace{q_0, \dots, q_0}_2, q_1, \dots, q_{N-n}) \right|_{q=q_0} \\ &= \dots = \frac{n}{n-1} \times \frac{n-1}{n-2} \times \dots \times \frac{m+1}{m} \times \left. \frac{d}{dq} E_{L,N}^{\text{I}}(\underbrace{q, \dots, q}_m, \underbrace{q_0, \dots, q_0}_{n-m}, q_1, \dots, q_{N-n}) \right|_{q=q_0} . \end{aligned}$$

The proof for the missed detection probability is the same. \square




Figure 4-6. A contour plot of the Bayes risk for two agents with the OR fusion rule varying q_1 and q_2 for $c_{10} = c_{01} = 1$, $p = 0.5$, and additive Gaussian noise with zero mean and unit variance. Arrow D_1 indicates the change of belief in the team of identical agents and arrow D_2 indicates that in the team of diverse agents. If there are errors in the perception, it is almost the worst case when the agents have identical errors.

Theorem 4.4. For any $m < n \leq N$ and any Beliefs q_0, q_1, \dots, q_{N-m} ,

$$\begin{aligned}
& \left. \frac{d}{dq} \left(c_{10} p_0 E_{L,N}^I(q, \dots, q, \underbrace{q_1, \dots, q_{N-n}}_n) + c_{01} p_1 E_{L,N}^{II}(q, \dots, q, \underbrace{q_1, \dots, q_{N-n}}_n) \right) \right|_{q=q_0} \\
&= \frac{n}{m} \frac{d}{dq} \left(c_{10} p_0 E_{L,N}^I(q, \dots, q, \underbrace{q_0, \dots, q_0}_m, \underbrace{q_1, \dots, q_{N-n}}_{n-m}) \right. \\
& \quad \left. + c_{01} p_1 E_{L,N}^{II}(q, \dots, q, \underbrace{q_0, \dots, q_0}_m, \underbrace{q_1, \dots, q_{N-n}}_{n-m}) \right) \Big|_{q=q_0}.
\end{aligned}$$

Proof. This theorem is directly obtained from Corollary 4.3. \square

Theorem 4.4 implies that if a team consists of n identical and $N - n$ diverse agents then the change of identical agents' belief has an impact on the change of the team Bayes risk proportional to n . Thus, as the team gets more diverse, the impact of change is decreased in a reciprocal fashion.

Figure 4-6 depicts the global Bayes risk for a team of two agents having various




Figure 4-7. Suppose that Alexis has the smallest threshold ($\lambda_{(1)}$) and Britta has the second smallest threshold ($\lambda_{(2)}$), etc. (a) If $Y < \lambda_{(1)}$, all agents will declare 0. (b) If $\lambda_{(1)} < Y < \lambda_{(2)}$, all agents except Alexis declare 0.

beliefs. A change of belief in the identical team is marked as the arrow D_1 . The increase of the Bayes risk is higher than the increase due to a change of belief in the diverse team, which is marked as the arrows D_2 . Even though Theorem 4.4 deals with the infinitesimal change of belief, Figure 4-6 also supports the similar trend for larger changes.

■ 4.2.2 A Team of Agents Who Observe the Same Signal

Now let us change the setting so that there is only one observation. This assumption is applicable to some scenarios, such as criminal trials in which all jurors observe the same evidence, such as the evidence presented by prosecutors or counsel.

In this case, it is meaningless to consider multiple identical agents. Since they have the same observation, their decisions will always coincide. On the other hand, for a diverse team, local decisions can be varied depending on the perceived belief of individual agents. For the same observation, it can happen that an agent whose belief is biased toward $H = 1$ chooses 1 while another agent whose belief is biased $H = 0$ chooses 0.

Even though the diverse agents have different beliefs, their decision making is not independent of each other. For example, if the common observation $Y = y$ is smaller than the minimum of the agents' decision thresholds as depicted in Figure 4-7a, then all agents will declare 0. If y is between the first and the second smallest decision thresholds as in Figure 4-7b, then the agent with the minimum decision threshold will declare 1 and everyone else will declare 0. It will not happen that the agent with the minimum decision threshold declares 0 and some others declare 1.

Theorem 4.5. *Under the L -out-of- N fusion rule, the final decision is solely determined by the agent who has the L th smallest belief on $H = 0$ among the N agents.*

Proof. Since decision threshold is monotonically increasing in the belief, this theorem can be rephrased as follows:

“Under the L -out-of- N fusion rule, the final decision is solely determined by the agent who has the L th smallest decision threshold among the N agents.”

Let $\lambda_{(L)}$ denote the L th smallest decision threshold. If $Y < \lambda_{(L)}$, at least $N - L + 1$ agents, whose decision thresholds are not less than $\lambda_{(L)}$, choose 0 and the global decision will also be 0. Otherwise, at least L agents, whose decision thresholds are not greater than $\lambda_{(L)}$, choose 1 and the global decision will be 1. Therefore, the global decision is always the same as the local decision made by the agent who has the L th smallest decision threshold. \square

The Bayes risk of the diverse team is simply given by

$$R_D = c_{10}p_0P_e^I(\lambda_{(L)}) + c_{01}p_1P_e^{II}(\lambda_{(L)}). \quad (4.10)$$

Please note that $\lambda_{(L)}$ should be determined by $q_{(L)}$ but not equal to $r_{L,N}(q_{(L)})$, where $q_{(L)}$ is the L th smallest numbers among q_1, \dots, q_n . The global Bayes risk (4.10) is computed not with global error probabilities P_E^I and P_E^{II} but with local error probabilities P_e^I and P_e^{II} . Thus the agents just need to perform decision making as if they are the only decision maker, i.e., $\lambda_{(L)} = r(q_{(L)})$, where $r(\cdot) = r_{1,1}(\cdot)$.

There is an analogy between this theorem and the median voter theorem. The median voter theorem states that the median voter determines the outcome of majority voting. In Theorem 4.5, the agent who has the median belief is pivotal for the majority fusion rule. The agent can be referred to as the median voter.

When diverse agents observe conditionally iid signals, they have an advantage of effectively higher signal-to-noise ratio than that of individual signals. It is why they can reduce their global Bayes risk even though they have different beliefs. However, when they observe the same signal, they do not have such an advantage. Their team

Bayes risk is lower bounded by $c_{10}p_0P_e^I(r(p_0)) + c_{01}p_1P_e^{II}(r(p_0))$ even though N goes to infinity. Instead, if their beliefs are distributed around the true prior probability p_0 , such as uniformly distributed within $[p_0 - \delta, p_0 + \delta]$ for some $\delta > 0$, then their median belief converges almost surely to p_0 and their Bayes risk will be almost the same as that of a single perfect agent.

■ 4.3 Experts and Novices in a Team

So far, teams of agents who observe conditionally iid signals have been considered. Then the perfect agents, who know p_0 , always outperform imperfect agents of the same number, who perceive p_0 as q_i .

In this section, agents observe conditionally independent but not identically distributed signals. Some agents have better signals with smaller noise and some others observe worse signals. These agents perform a decision making as a team but do not know how good others' observations are.

The sets of perfect agents and imperfect agents are identical with respect to the likelihoods of their observations, e.g., both perfect Agent n and imperfect Agent n observe signals Y_n whose likelihoods are equally $f_{Y_n|H}(y_n|h)$. They are different only in that perfect Agent n knows what p_0 is but imperfect Agent n does not.

Even though perfect agents know the likelihoods of the hypotheses and their own observations, they do not know likelihoods of others' observations. This ignorance about other agents is the cause of their suboptimality. Intuitively, their decision-making rules can be optimized only if they have all information, including the likelihoods of others' observations. They do not, however, and cannot figure out the optimal rule. On the other hand, some imperfect agents can have the optimal rules if their wrong perceived beliefs luckily lead to the optimal rules. Ironically, two errors may cancel out each other but one error cannot cancel out itself.

Before discussing the performance of perfect and imperfect agents, let us think about the optimal decision-making rule. Bayes risk has the same form as (4.7). For example, for $N = 2$ and the 1-out-of-2 fusion rule, the agents' optimal decision

thresholds λ_1^* and λ_2^* are the solution to the following two equations:

$$\begin{aligned}\frac{f_{Y_1|H}(\lambda_1|1)}{f_{Y_1|H}(\lambda_1|0)} &= \frac{c_{10}p_0(1 - P_{e,2}^I(\lambda_2))}{c_{01}p_1P_{e,2}^{II}(\lambda_2)} = \frac{c_{10}p_0 \int_{\lambda_2}^{\infty} f_{Y_2|H}(y_2|0) dy_2}{c_{01}p_1 \int_{\lambda_2}^{\infty} f_{Y_2|H}(y_2|1) dy_2}, \\ \frac{f_{Y_2|H}(\lambda_2|1)}{f_{Y_2|H}(\lambda_2|0)} &= \frac{c_{10}p_0 \int_{\lambda_1}^{\infty} f_{Y_1|H}(y_1|0) dy_1}{c_{01}p_1 \int_{\lambda_1}^{\infty} f_{Y_1|H}(y_1|1) dy_1}.\end{aligned}\tag{4.11}$$

These equations look identical to (4.5) but $\lambda_1 \neq r_{L,N}(p_0)$ because $f_{Y_1|H}(y|h) \neq f_{Y_2|H}(y|h)$ for any y and h . Each agent requires both $f_{Y_1|H}$ and $f_{Y_2|H}$ to solve them.

However, they only know the likelihood functions of their own observations. In the perfect-agent case, when agents know the prior probability p_0 , the following pair of equations are the modified version of (4.11) from Agent 1's perspective:

$$\begin{aligned}\frac{f_{Y_1|H}(\lambda_1|1)}{f_{Y_1|H}(\lambda_1|0)} &= \frac{c_{10}p_0 \int_{\lambda_2}^{\infty} f_{Y_1|H}(y_2|0) dy_2}{c_{01}p_1 \int_{\lambda_2}^{\infty} f_{Y_1|H}(y_2|1) dy_2}, \\ \frac{f_{Y_1|H}(\lambda_2|1)}{f_{Y_1|H}(\lambda_2|0)} &= \frac{c_{10}p_0 \int_{\lambda_1}^{\infty} f_{Y_1|H}(y_1|0) dy_1}{c_{01}p_1 \int_{\lambda_1}^{\infty} f_{Y_1|H}(y_1|1) dy_1}.\end{aligned}\tag{4.12}$$

Agent 1 will compute the solution $\lambda_1^{(1)}$ and $\lambda_2^{(1)}$ of (4.12); she takes $\lambda_1^{(1)}$ as her decision threshold and assumes $\lambda_2^{(1)}$ as Agent 2's threshold.

Likewise, perfect Agent 2 develops the following equations and finds their solution $\lambda_1^{(2)}$ and $\lambda_2^{(2)}$:

$$\begin{aligned}\frac{f_{Y_2|H}(\lambda_1|1)}{f_{Y_2|H}(\lambda_1|0)} &= \frac{c_{10}p_0 \int_{\lambda_2}^{\infty} f_{Y_2|H}(y_2|0) dy_2}{c_{01}p_1 \int_{\lambda_2}^{\infty} f_{Y_2|H}(y_2|1) dy_2}, \\ \frac{f_{Y_2|H}(\lambda_2|1)}{f_{Y_2|H}(\lambda_2|0)} &= \frac{c_{10}p_0 \int_{\lambda_1}^{\infty} f_{Y_2|H}(y_1|0) dy_1}{c_{01}p_1 \int_{\lambda_1}^{\infty} f_{Y_2|H}(y_1|1) dy_1}.\end{aligned}\tag{4.13}$$

She chooses her decision threshold as $\lambda_2^{(2)}$. The decision rules $\lambda_1^{(1)}$ determined by (4.12) and $\lambda_2^{(2)}$ by (4.13) can be said to be rational but not optimal because they are different from the optimal ones λ_1^* and λ_2^* determined by (4.11).

In the imperfect-agent case, agents use their perceived beliefs in the decision making. Imperfect Agents 1 and 2 respectively design their decision rules from the

following equations, which are differentiated from (4.12) and (4.13).

$$\begin{aligned}
(\text{Imperfect Agent 1}) \quad & \frac{f_{Y_1|H}(\lambda_1|1)}{f_{Y_1|H}(\lambda_1|0)} = \frac{c_{10}q_1 \int_{\lambda_2}^{\infty} f_{Y_1|H}(y_2|0) dy_2}{c_{01}(1-q_1) \int_{\lambda_2}^{\infty} f_{Y_1|H}(y_2|1) dy_2}, \\
& \frac{f_{Y_1|H}(\lambda_2|1)}{f_{Y_1|H}(\lambda_2|0)} = \frac{c_{10}q_1 \int_{\lambda_1}^{\infty} f_{Y_1|H}(y_1|0) dy_1}{c_{01}(1-q_1) \int_{\lambda_1}^{\infty} f_{Y_1|H}(y_1|1) dy_1}. \quad (4.14)
\end{aligned}$$

$$\begin{aligned}
(\text{Imperfect Agent 2}) \quad & \frac{f_{Y_2|H}(\lambda_1|1)}{f_{Y_2|H}(\lambda_1|0)} = \frac{c_{10}q_2 \int_{\lambda_2}^{\infty} f_{Y_2|H}(y_2|0) dy_2}{c_{01}(1-q_2) \int_{\lambda_2}^{\infty} f_{Y_2|H}(y_2|1) dy_2}, \\
& \frac{f_{Y_2|H}(\lambda_2|1)}{f_{Y_2|H}(\lambda_2|0)} = \frac{c_{10}q_2 \int_{\lambda_1}^{\infty} f_{Y_2|H}(y_1|0) dy_1}{c_{01}(1-q_2) \int_{\lambda_1}^{\infty} f_{Y_2|H}(y_1|1) dy_1}. \quad (4.15)
\end{aligned}$$

Note that these equations imply that each agent acts as if the other agent's perceived belief is the same as hers.

The performance of the imperfect agents depends on their perceived beliefs: it can be better or worse than that of the perfect agents. For given likelihood functions $f_{Y_1|H}$ and $f_{Y_2|H}$, the rational decision rules of the imperfect agents are arbitrary according to q_1 and q_2 while the rational decision rules of the perfect agents are fixed. Ideally, it will be the optimal situation for the imperfect agents if their beliefs accidentally lead to the optimal decision thresholds λ_1^* and λ_2^* . In other words, if Agent 1 has q_1 and Agent 2 has q_2 such that the solution of (4.14) is equal to λ_1^* and the solution of (4.15) is equal to λ_2^* , then their Bayes risk will be minimal and lower than that of the perfect agents. However, it is also possible that the imperfect agents have beliefs such that their decision rules are worse than $\lambda_1^{(1)}$ and $\lambda_2^{(2)}$.

For example, suppose that the agents observe signals $Y_n = H + W_n$, where W_n are zero-mean Gaussian random variables. Agent 1 observes a better signal with $W_1 \sim \mathcal{N}(0, 0.5)$ and Agent 2 observes a worse signal with $W_2 \sim \mathcal{N}(0, 1)$. For convenience, let's call Agent 1 an *expert* and Agent 2 a *novice*.

Figure 4-8 shows Bayes risk for various perceived beliefs for a case when the agents make decision with the 1-out-of-2 fusion rule and $p_0 = 0.3$. The region within the dashed boundary is the set of beliefs that lead to better performance than per-




Figure 4-8. Team Bayes risk for various beliefs. Agent 1 (expert) observes $Y_1 = H + W_1$, where $W_1 \sim \mathcal{N}(0, 0.5)$, and Agent 2 (novice) observes $Y_2 = H + W_2$, where $W_2 \sim \mathcal{N}(0, 1)$. The fusion rule is the OR rule; $p_0 = 0.3$ and $c_{10} = c_{01} = 1$.

fect agents. In the optimal point, the expert slightly underestimates and the novice overestimates the likelihood of $H = 0$. This pattern holds for any prior probability; the optimality is achieved by a team of underestimating expert and overestimating novice, Figure 4-9a.

There is not an all-round pattern that always holds. For the 2-out-of-2 fusion rule, a team of overestimating expert and underestimating novice performs better, as depicted in Figure 4-9b.

Even though there is not a strict pattern for optimal beliefs, it is always true that the decision making by a team of perfect agents is suboptimal if the agents do not observe equally noisy signals. For imperfect agents, their performance depends on their perceived beliefs. They can outperform perfect agents and moreover achieve the optimal decision making if their team consists of properly imperfect agents.

■ 4.4 Conclusion

When agents do not know the prior probability, they can form a team of agents who make a decision together to improve their decision making. This chapter models

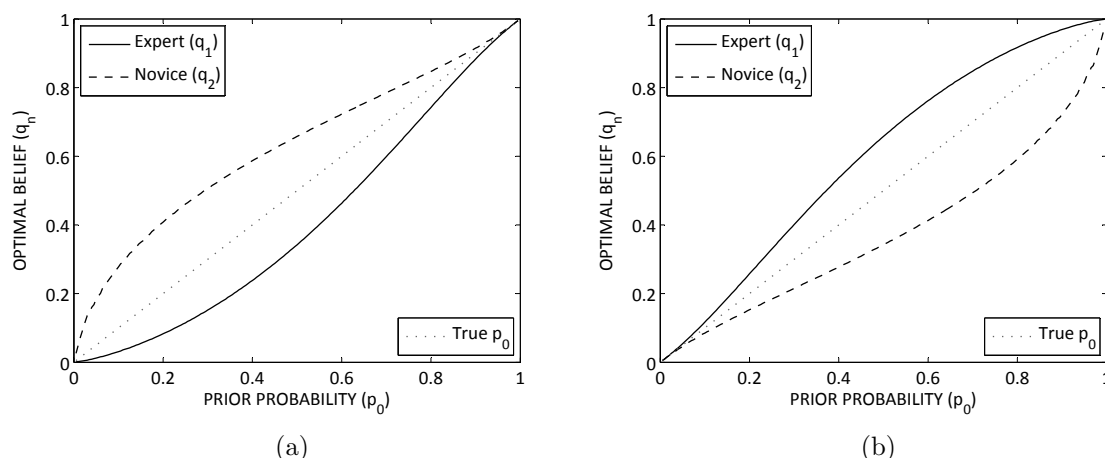


Figure 4-9. Beliefs q_1 and q_2 that minimize the team Bayes risk for various p_0 . The expert observes $Y_1 = H + W_1$, where $W_1 \sim \mathcal{N}(0, 0.5)$, and the novice observes $Y_2 = H + W_2$, where $W_2 \sim \mathcal{N}(0, 1)$. Having these wrong beliefs yield better performance than knowing the correct prior probability. (a) The OR fusion rule. (b) The AND fusion rule.

the team as a distributed detection system with a fixed symmetric fusion rule. The performance of a team of identical agents, who equally perceive the probability, is compared to that of diverse agents, who differently perceive it. Which team is better depends on their beliefs and the prior probability. However, the team of diverse agents is more desirable from a stability perspective; their team Bayes risk does not change as much as that of the team of identical agents when the agents change their beliefs.

A team of imperfect agents cannot outperform perfect agents of the same number if the agents observe conditionally iid signals. However, the former can perform even better than the latter otherwise. Since the perfect agents know the prior probability but not the likelihood functions of other agents' signals, their rational decision rules are always suboptimal. The imperfect agents, on the other hand, do not know both the prior probability and the likelihood functions of other agents. Therefore, it can happen that their wrong beliefs accidentally compensate their misunderstanding of the likelihood functions of other agents. Then they will have the optimal decision rules.

■ 4.A Proof of Theorem 4.1

The beliefs q_0, \dots, q_{N-n} are constant while q can be varied. The function $E_{L,N}^I(q, \dots, q, q_1, \dots, q_{N-n})$ can be split into variable and constant terms:

$$\begin{aligned}
& E_{L,N}^I(\underbrace{q, \dots, q}_n, q_1, \dots, q_{N-n}) \\
&= \sum_{\ell=\max\{0, L-N+n\}}^n [\mathbb{P}\{\ell \text{ among the first } n \text{ agents raise false alarms}\} \\
&\quad \times \mathbb{P}\{L-\ell \text{ or more among the next } N-n \text{ agents raise false alarms}\}] \\
&= \sum_{\ell=\max\{0, L-N+n\}}^n \left[\binom{n}{\ell} (e^I(q))^\ell (1-e^I(q))^{n-\ell} \sum_{m=\max\{0, L-\ell\}}^{N-n} FA_{N-n}^m \right],
\end{aligned}$$

where $e^I(q) = P_e^I(r_{L,N}(q))$ denotes the local false alarm probability when the agent has a belief q and FA_{N-n}^m denotes the probability that exactly m agents among the last $N-n$ agents who respectively have a fixed belief q_1, q_2, \dots, q_{N-n} raise false alarm. The terms FA_{N-n}^m are constant.

The first derivative of the variable terms is

$$\begin{aligned}
& \left. \frac{d}{dq} \left(\binom{n}{\ell} (e^I(q))^\ell (1-e^I(q))^{n-\ell} \right) \right|_{q=q_0} \\
&= \left(\ell \binom{n}{\ell} (e^I(q_0))^{\ell-1} (1-e^I(q_0))^{n-\ell} - (n-\ell) \binom{n}{\ell} (e^I(q_0))^\ell (1-e^I(q_0))^{n-\ell-1} \right) \dot{e}^I(q_0) \\
&\stackrel{(a)}{=} n \left(\binom{n-1}{\ell-1} (e^I(q_0))^{\ell-1} (1-e^I(q_0))^{n-\ell} - \binom{n-1}{\ell} (e^I(q_0))^\ell (1-e^I(q_0))^{n-\ell-1} \right) \dot{e}^I(q_0),
\end{aligned}$$

where $\dot{e}^I(q_0) \triangleq \left. \frac{de^I(q)}{dq} \right|_{q=q_0}$. The equality (a) holds because

$$\ell \binom{n}{\ell} = n \binom{n-1}{\ell-1} \quad \text{and} \quad (n-\ell) \binom{n}{\ell} = n \binom{n-1}{\ell}.$$

Thus, the first derivative of $E_{L,N}^I(q, \dots, q, q_1, \dots, q_{N-n})$ at $q = q_0$ is

$$\begin{aligned}
& \left. \frac{d}{dq} E_{L,N}^I(q, \dots, q, q_1, \dots, q_{N-n}) \right|_{q=q_0} \\
&= n \sum_{\ell=\max\{1, L-N+n\}}^n \left[\binom{n-1}{\ell-1} (e^I(q_0))^{\ell-1} (1-e^I(q_0))^{n-\ell} \dot{e}^I(q_0) \sum_{m=\max\{0, L-\ell\}}^{N-n} FA_{N-n}^m \right] \\
&\quad - n \sum_{\ell=\max\{0, L-N+n\}}^{n-1} \left[\binom{n-1}{\ell} (e^I(q_0))^\ell (1-e^I(q_0))^{n-\ell-1} \dot{e}^I(q_0) \sum_{m=\max\{0, L-\ell\}}^{N-n} FA_{N-n}^m \right] \\
&= n \sum_{\ell=\max\{0, L-N+n-1\}}^{n-1} \left[\binom{n-1}{\ell} (e^I(q_0))^\ell (1-e^I(q_0))^{n-\ell-1} \dot{e}^I(q_0) \sum_{m=\max\{0, L-\ell-1\}}^{N-n} FA_{N-n}^m \right] \\
&\quad - n \sum_{\ell=\max\{0, L-N+n\}}^{n-1} \left[\binom{n-1}{\ell} (e^I(q_0))^\ell (1-e^I(q_0))^{n-\ell-1} \dot{e}^I(q_0) \sum_{m=\max\{0, L-\ell\}}^{N-n} FA_{N-n}^m \right],
\end{aligned}$$

where some of the terms in the first summation and the second summation are canceled out. The leftover terms are

$$\begin{aligned}
& \left. \frac{d}{dq} E_{L,N}^I(q, \dots, q, q_1, \dots, q_{N-n}) \right|_{q=q_0} \\
&= n \sum_{\ell=\max\{0, L-N+n\}}^{n-1} \left[\binom{n-1}{\ell} (e^I(q_0))^\ell (1-e^I(q_0))^{n-\ell-1} \dot{e}^I(q_0) FA_{N-n}^{L-\ell-1} \right] \\
&\quad + n \binom{n-1}{L-N+n-1} (e^I(q_0))^{L-N+n-1} (1-e^I(q_0))^{N-L} \dot{e}^I(q_0) \sum_{m=\max\{0, N-n\}}^{N-n} FA_{N-n}^m.
\end{aligned}$$

The last term can be combined into the summation because

$$\sum_{m=\max\{0, N-n\}}^{N-n} FA_{N-n}^m = FA_{N-n}^{N-n}.$$

Therefore,

$$\begin{aligned}
& \left. \frac{d}{dq} E_{L,N}^I(q, \dots, q, q_1, \dots, q_{N-n}) \right|_{q=q_0} \\
&= n \sum_{\ell=\max\{0, L-N+n-1\}}^{n-1} \left[\binom{n-1}{\ell} (e^I(q_0))^\ell (1-e^I(q_0))^{n-\ell-1} \dot{e}^I(q_0) FA_{N-n}^{L-\ell-1} \right] \\
&= n \dot{e}^I(q_0) X_n,
\end{aligned} \tag{4.16}$$

where

$$X_n \triangleq \sum_{\ell=\max\{0, L-N+n-1\}}^{n-1} \left[\binom{n-1}{\ell} (e^I(q_0))^\ell (1 - e^I(q_0))^{n-\ell-1} FA_{N-n}^{L-\ell-1} \right].$$

Next, let us consider the global false alarm probability $E_{L,N}^I(q, \dots, q, q_0, \dots, q_{N-n})$ when the first $n-1$ agents change their belief while the other agents have beliefs q_0, q_1, \dots, q_N . Its first derivative at $q = q_0$ can be computed similarly:

$$\begin{aligned} & \left. \frac{d}{dq} E_{L,N}^I(q, \dots, q, q_0, q_1, \dots, q_{N-n}) \right|_{q=q_0} \\ &= (n-1) \sum_{\ell=\max\{0, L-N+n-2\}}^{n-2} \left[\binom{n-2}{\ell} (e^I(q_0))^\ell (1 - e^I(q_0))^{n-\ell-2} \dot{e}^I(q_0) FA_{N-n+1}^{L-\ell-1} \right] \\ &= (n-1) \dot{e}^I(q_0) X_{n-1}, \end{aligned} \tag{4.17}$$

where FA_{N-n+1}^m denotes the probability that m agents among the last $N-n+1$ agents who respectively have a fixed belief q_0, q_1, \dots, q_{N-n} raise false alarm, and

$$X_{n-1} \triangleq \sum_{\ell=\max\{0, L-N+n-2\}}^{n-2} \left[\binom{n-2}{\ell} (e^I(q_0))^\ell (1 - e^I(q_0))^{n-\ell-2} FA_{N-n+1}^{L-\ell-1} \right].$$

The physical meaning of X_n and X_{n-1} implies that they are the same. First,

$$\begin{aligned} X_n &= \sum_{\ell=\max\{0, L-N+n-1\}}^{n-1} \left[\binom{n-1}{\ell} (e^I(q_0))^\ell (1 - e^I(q_0))^{n-\ell-1} FA_{N-n}^{L-\ell-1} \right] \\ &= \sum_{\ell=\max\{0, L-N+n-1\}}^{n-1} \left[\mathbb{P}\{\ell \text{ among the first } n-1 \text{ agents raise false alarms}\} \right. \\ &\quad \left. \times \mathbb{P}\{L-\ell-1 \text{ among the remaining } N-n \text{ agents raise false alarms}\} \right] \\ &= \mathbb{P}\{L-1 \text{ among the } N-1 \text{ agents raise false alarms}\}. \end{aligned}$$

Note that, among the $N-1$ agents, the first $n-1$ agents identically have beliefs q_0 and the next $N-n$ agents respectively have beliefs q_1, \dots, q_{N-n} .

Second,

$$\begin{aligned}
X_{n-1} &= \sum_{\ell=\max\{0, L-N+n-2\}}^{n-2} \left[\binom{n-2}{\ell} (e^I(q_0))^\ell (1 - e^I(q_0))^{n-\ell-2} FA_{N-n+1}^{L-\ell-1} \right] \\
&= \sum_{\ell=\max\{0, L-N+n-2\}}^{n-2} [\mathbb{P}\{\ell \text{ among the first } n-2 \text{ agents raise false alarms}\}] \\
&\quad \times \mathbb{P}\{L-\ell-1 \text{ among the remaining } N-n+1 \text{ agents raise false alarms}\}] \\
&= \mathbb{P}\{L-1 \text{ among the } N-1 \text{ agents raise false alarms}\}.
\end{aligned}$$

Again, among the $N-1$ agents, the first $n-2$ agents identically have beliefs q_0 and the next $N-n+1$ agents respectively have beliefs q_0, q_1, \dots, q_{N-n} . Therefore, X_n and X_{n-1} are the probabilities of identical events and are essentially equal.

The theorem is proven from (4.16) and (4.17) with $X_n = X_{n-1}$.

Sequential Detection with Social Learning

Another simple decision-making structure for multiple agents is a sequential decision-making. Consider multiple agents performing the same binary hypothesis testing. Unlike the agents within a decision-making team in Chapter 4, these agents are individually responsible for their own decisions. These agents, who only care about maximizing their own payoffs, i.e., minimizing their own Bayes risks, are often called *selfish* agents.

Each agent observes a signal, which is not visible to other agents, to make a decision. This signal is called a *private signal*. Suppose that the agents make decisions sequentially in a predetermined order and the decisions are visible to other agents, Figure 5-1. The decisions that are publicly observed are called *public signals* to be distinguished from private signals.

The public signals make earlier-acting agents and later-acting agents different. Later-acting agents can learn some information about the right (or better) choice from earlier-acting agents' decisions and reflect it in their own choices. The learning behavior is called *social learning*.

The agents are named Alexis, Britta, Carol, etc., according to their decision-making order. As each agent makes a decision, her decision is shown to the other agents. Our interest is the decision-making of the last agent, Norah. As she observes precedent decisions, she updates her belief to elaborate her likelihood ratio test. The update process depends on her initial belief and the history of decisions. However,




Figure 5-1. A sequential decision making model with N agents (Alexis, Britta, \dots , Norah). Agent n can observe $n - 1$ decisions made by all precedent agents.

it will be different from what was explained in Section 2.3 because the agents have incorrect individual beliefs.

It turns out that the incorrect beliefs can result in smaller Bayes risk for Norah than correct beliefs do. In addition, when the private signals are generated from iid Gaussian likelihood functions, the optimal beliefs follow a systematic pattern: the earlier-acting agents should act as if the prior probability of the unlikely hypothesis is larger than it is in reality, and vice versa. This is interpreted as being *open minded* toward the unlikely hypothesis.

Section 5.1 explains agents' belief updates and decision-making rule, position-by-position. Section 5.2 discusses the beliefs that minimize the Bayes risk of the last agent. It provides the proof that the correct belief is not the minimizer and shows that there is a systematic pattern of the minimizer. Section 5.3 concludes the chapter.

■ 5.1 Social Learning of Imperfect Agents

Throughout this chapter, the agents in Figure 5-1 are assumed to have different beliefs q_1, q_2, \dots, q_N , respectively. Agent n observes a conditionally iid private signal $Y_n = H + W_n$, where $W_n \sim \mathcal{N}(0, \sigma^2)$:

$$f_{Y_n|H}(y_n | 0) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{y_n^2}{2\sigma^2}\right], \quad (5.1a)$$

$$f_{Y_n|H}(y_n | 1) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(y_n - 1)^2}{2\sigma^2}\right]. \quad (5.1b)$$

Each agent has the costs of errors c_{10} and c_{01} that are the same as other agents' costs of errors but only the outcome of her own decisions matters to her. In other words, the Bayes risk of Agent n is determined by her own false alarm and missed detection probabilities:

$$R_n = c_{10}p_0\mathbb{P}\{\widehat{H}_n = 1 \mid H = 0\} + c_{01}(1 - p_0)\mathbb{P}\{\widehat{H}_n = 0 \mid H = 1\}.$$

Later-acting agents utilize public signals as well as their private signals for Bayesian hypothesis testing. They can do perfect social learning if all agents know the prior probability p_0 . If they do not, then at least they need to know other agents' beliefs so that they can understand on what belief the decisions are made. However, agents in this model know neither p_0 nor other agents' beliefs q_i . Therefore their social learning is inevitably incomplete. Their social learning is explained with a position-wise decision-making strategy.

In this chapter and Chapter 6, in which social learning is considered, the following notations are used. Superscript characters, such as A, B, etc., mean “upon observing Alexis’s decision, Britta’s decision, etc.”; to specify their decision values, 0 or 1 is used instead of the Roman alphabet. For example, q_3^{AB} denotes the (updated) belief of Agent 3, Carol, upon observing Alexis’s and Britta’s decisions \widehat{H}_1 and \widehat{H}_2 , and q_3^{01} denotes Carol’s updated belief upon observing $\widehat{H}_1 = 0$ and $\widehat{H}_2 = 1$. Subscript letters, such as A, B, etc., mean “that Alexis thinks,” “that Britta thinks,” etc. For example, Britta thinks that the probability of Alexis choosing 0 when the true state is 0 is $p_{\widehat{H}_1|H}(0|0)_B$. We need to clarify it because the agents are not aware of others’ prior beliefs so they will misunderstand public signals. The meaning of the subscript will be explained in detail in Section 5.1.2.

■ 5.1.1 Alexis, the First Agent

Alexis only observes a private signal Y_1 . She uses the following likelihood ratio test with Y_1 and her belief q_1 :

$$\frac{f_{Y_1|H}(y_1|1)}{f_{Y_1|H}(y_1|0)} \underset{\widehat{H}_1(y_1)=0}{\overset{\widehat{H}_1(y_1)=1}{\gtrless}} \frac{c_{10}q_1}{c_{01}(1-q_1)}. \quad (5.2)$$

Since the likelihood ratio is increasing in y_1 , the likelihood ratio test can be simplified to comparison with an appropriate decision threshold:

$$y_1 \underset{\widehat{H}_1(y_1)=0}{\overset{\widehat{H}_1(y_1)=1}{\gtrless}} r(q_1), \quad (5.3)$$

where $r(q)$ denotes the decision threshold λ that satisfies

$$\frac{f_{Y|H}(\lambda|1)}{f_{Y|H}(\lambda|0)} = \frac{c_{10}q}{c_{01}(1-q)}. \quad (5.4)$$

■ 5.1.2 Britta, the Second Agent

Britta observes Alexis's decision \widehat{H}_1 as well as a private signal Y_2 . She updates her belief from q_2 to q_2^A :

$$\begin{aligned} q_2^A &= p_{H|\widehat{H}_1}(0|\widehat{h}_1)_B \\ &= \frac{q_2 p_{\widehat{H}_1|H}(\widehat{h}_1|0)_B}{q_2 p_{\widehat{H}_1|H}(\widehat{h}_1|0)_B + (1-q_2) p_{\widehat{H}_1|H}(\widehat{h}_1|1)_B}. \end{aligned} \quad (5.5)$$

Note that Britta needs q_1 to compute the conditional probabilities

$$\begin{aligned} p_{\widehat{H}_1|H}(0|h) &= \mathbb{P}(\{Y_1 < r(q_1) | H = h\}) = \int_{-\infty}^{r(q_1)} f_{Y|H}(y|h) dy, \\ p_{\widehat{H}_1|H}(1|h) &= \mathbb{P}(\{Y_1 \geq r(q_1) | H = h\}) = \int_{r(q_1)}^{\infty} f_{Y|H}(y|h) dy. \end{aligned}$$

Since she does not know q_1 , however, she treats \widehat{H}_1 as if Alexis has belief q_2 , the same as hers:

$$p_{\widehat{H}_1|H}(0|h)_B = \mathbb{P}(\{Y_1 < r(q_2) | H = h\}) = \int_{-\infty}^{r(q_2)} f_{Y|H}(y|h) dy, \quad (5.6a)$$

$$p_{\widehat{H}_1|H}(1|h)_B = \mathbb{P}(\{Y_1 \geq r(q_2) | H = h\}) = \int_{r(q_2)}^{\infty} f_{Y|H}(y|h) dy. \quad (5.6b)$$

The subscripts “B” in (5.5) and (5.6) indicate that the probabilities are computed based on Britta’s belief q_2 . Thus, these probabilities are not computed correctly.

An interesting observation is that Alexis’s biased belief q_1 does not affect Britta’s belief update. There is no trace of q_1 in (5.5) and (5.6). Suppose that Alexis knows true prior probability p_0 and uses the decision threshold $r(p_0)$. Still Britta, who does not know what belief Alexis has, thinks that the conditional probability of Alexis choosing $\widehat{H}_1 = 0$ is given by (5.6) and updates her belief as in (5.5). It is clear in (5.5) that the updated belief depends only on Britta’s belief and Alexis’s decision.

However, Alexis’s prior belief still affects Britta’s performance with respect to her error probabilities. Alexis’s biased belief changes the probability of her decision. The changed probability is embedded in the probability of Britta’s decision:¹

$$\begin{aligned} p_{\widehat{H}_2|H}(\widehat{h}_2|0) &= \sum_{\widehat{h}_1} p_{\widehat{H}_2, \widehat{H}_1|H}(\widehat{h}_2, \widehat{h}_1|0) \\ &= p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|0,0)_B p_{\widehat{H}_1|H}(0|0)_A + p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|1,0)_B p_{\widehat{H}_1|H}(1|0)_A, \end{aligned} \quad (5.7a)$$

$$\begin{aligned} p_{\widehat{H}_2|H}(\widehat{h}_2|1) &= \sum_{\widehat{h}_1} p_{\widehat{H}_2, \widehat{H}_1|H}(\widehat{h}_2, \widehat{h}_1|1) \\ &= p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|0,1)_B p_{\widehat{H}_1|H}(0|1)_A + p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|1,1)_B p_{\widehat{H}_1|H}(1|1)_A. \end{aligned} \quad (5.7b)$$

In conclusion, Alexis’s belief changes Britta’s error probabilities but not her decision rule.

¹The subscripts “A” are written in (5.7) for clarification that $p_{\widehat{H}_1|H}(\cdot|\cdot)$ computed based on Alexis’s belief q_1 should be used unlike in (5.5).

■ 5.1.3 Carol, the Third Agent

Carol updates her belief upon observing public signals \widehat{H}_1 and \widehat{H}_2 :

$$\begin{aligned} q_3^{\text{AB}} &= p_{H|\widehat{H}_1, \widehat{H}_2}(0|\widehat{h}_1, \widehat{h}_2)_C \\ &= \frac{q_3 p_{\widehat{H}_1, \widehat{H}_2|H}(\widehat{h}_1, \widehat{h}_2|0)_C}{q_3 p_{\widehat{H}_1, \widehat{H}_2|H}(\widehat{h}_1, \widehat{h}_2|0)_C + (1 - q_3) p_{\widehat{H}_1, \widehat{H}_2|H}(\widehat{h}_1, \widehat{h}_2|1)_C}. \end{aligned}$$

Note that Carol uses her belief q_3 to compute the conditional probabilities as subscripts “C” indicate.

The belief update process can also be expressed in a cascade form:

$$\begin{aligned} \frac{q_3^{\text{AB}}}{1 - q_3^{\text{AB}}} &= \frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_2, \widehat{H}_1|H}(\widehat{h}_2, \widehat{h}_1|0)_C}{p_{\widehat{H}_2, \widehat{H}_1|H}(\widehat{h}_2, \widehat{h}_1|1)_C} \\ &= \left(\frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)_C}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)_C} \right) \frac{p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|\widehat{h}_1, 0)_C}{p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|\widehat{h}_1, 1)_C}. \end{aligned} \quad (5.8)$$

This update process can be split into two steps. The first step is to update her initial belief based on Alexis’s decision:

$$\frac{q_3^{\text{A}}}{1 - q_3^{\text{A}}} = \frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)_C}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)_C}. \quad (5.9)$$

The second step is to update her belief from q_3^{A} based on Britta’s decision:

$$\frac{q_3^{\text{AB}}}{1 - q_3^{\text{AB}}} = \frac{q_3^{\text{A}}}{1 - q_3^{\text{A}}} \frac{p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|\widehat{h}_1, 0)_C}{p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|\widehat{h}_1, 1)_C}. \quad (5.10)$$

Details of computations of (5.9) and (5.10) are as follows:

$$p_{\widehat{H}_1|H}(0|h)_C = \int_{-\infty}^{r(q_3)} f_{Y_1|H}(y|h) dy, \quad (5.11a)$$

$$p_{\widehat{H}_1|H}(1|h)_C = \int_{r(q_3)}^{\infty} f_{Y_1|H}(y|h) dy. \quad (5.11b)$$

Substituting (5.11) in (5.9), Carol can compute q_3^{A} for $\widehat{H}_1 = 0$ and $\widehat{H}_1 = 1$ respectively:

$$q_3^0 = \frac{q_3}{q_3 + (1 - q_3) \frac{\int_{-\infty}^{r(q_3)} f_{Y_1|H}(y|1) dy}{\int_{-\infty}^{r(q_3)} f_{Y_1|H}(y|0) dy}}, \quad (5.12a)$$

$$q_3^1 = \frac{q_3}{q_3 + (1 - q_3) \frac{\int_{r(q_3)}^{\infty} f_{Y_1|H}(y|1) dy}{\int_{r(q_3)}^{\infty} f_{Y_1|H}(y|0) dy}}. \quad (5.12b)$$

Then,

$$p_{\widehat{H}_2|\widehat{H}_1,H}(0|\widehat{h}_1, h)_3 = \mathbb{P}(\{Y_2 < r(q_3^A) | H = h\}) = \int_{-\infty}^{r(q_3^A)} f_{Y_2|H}(y|h) dy, \quad (5.13a)$$

$$p_{\widehat{H}_2|\widehat{H}_1,H}(1|\widehat{h}_1, h)_3 = \mathbb{P}(\{Y_2 \geq r(q_3^A) | H = h\}) = \int_{r(q_3^A)}^{\infty} f_{Y_2|H}(y|h) dy. \quad (5.13b)$$

Even though the value of \widehat{h}_1 may not seem to be used in (5.13), it is inherent in q_3^A and affects the computation results. Carol's updated belief q_3^{AB} is obtained by substituting (5.12) and (5.13) in (5.10).

The cascade form (5.8) suggests a recurrence relation of belief update. It will be discussed in Section 5.1.4

■ 5.1.4 Norah, the N th Agent

Norah, the N th agent, observes Y_N and $\widehat{H}_1, \dots, \widehat{H}_{N-1}$. Her belief update process mimics (5.8):

$$\frac{q_N^{AB \dots M}}{1 - q_N^{AB \dots M}} = \frac{q_N}{1 - q_N} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)_N}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)_N} \times \dots \times \frac{p_{\widehat{H}_{N-1}|\widehat{H}_1, \dots, \widehat{H}_{N-2}, H}(\widehat{h}_{N-1}|\widehat{h}_1, \dots, \widehat{h}_{N-2}, 0)_N}{p_{\widehat{H}_{N-1}|\widehat{H}_1, \dots, \widehat{H}_{N-2}, H}(\widehat{h}_{N-1}|\widehat{h}_1, \dots, \widehat{h}_{N-2}, 1)_N}.$$

As in Carol's case, she can update her belief step by step for each public signal. Let $U_n : [0, 1] \times \{0, 1\}^{n-1} \mapsto [0, 1]$ denote a general prior belief update function upon observing decisions of the first $n-1$ agents. The function U_n has the following recurrence relation:

- For $n = 1$, $U_1(q) = q$.

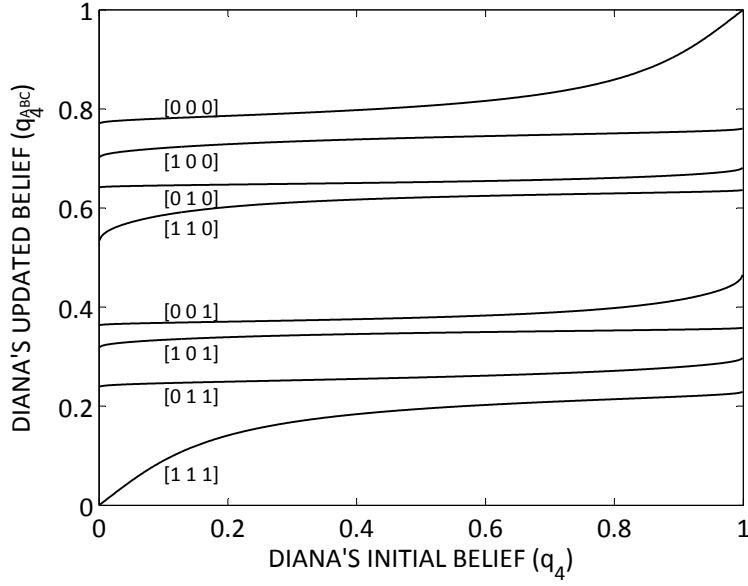


Figure 5-2. The function $U_4(q_4, \widehat{h}_1, \widehat{h}_2, \widehat{h}_3)$ —updated belief of Diana, the fourth agent (q_4^{ABC})—for each possible combination of Alexis’s, Britta’s, and Carol’s decisions $[\widehat{h}_1 \widehat{h}_2 \widehat{h}_3]$ when $c_{10} = c_{01} = 1$ and $\sigma^2 = 1$. Her updated belief is mostly dependent on Carol’s decision; the top four curves are for $\widehat{h}_3 = 0$ and the bottom four curves are for $\widehat{h}_3 = 1$.

- For $n > 1$,

$$U_n(q, \widehat{h}_1, \dots, \widehat{h}_{n-2}, 0) = \frac{\tilde{q}}{\tilde{q} + (1 - \tilde{q}) \frac{\int_{-\infty}^{r(\tilde{q})} f_{Y_{n-1}|H}(y|1) dy}{\int_{-\infty}^{r(\tilde{q})} f_{Y_{n-1}|H}(y|0) dy}},$$

$$U_n(q, \widehat{h}_1, \dots, \widehat{h}_{n-2}, 1) = \frac{\tilde{q}}{\tilde{q} + (1 - \tilde{q}) \frac{\int_{\lambda(\tilde{q})}^{\infty} f_{Y_{n-1}|H}(y|1) dy}{\int_{\lambda(\tilde{q})}^{\infty} f_{Y_{n-1}|H}(y|0) dy}},$$

where $\tilde{q} = U_{n-1}(q, \widehat{h}_1, \dots, \widehat{h}_{n-2})$.

Norah’s final updated belief is

$$q_N^{\text{AB}\dots\text{M}} = U_N(q_N, \widehat{h}_1, \widehat{h}_2, \dots, \widehat{h}_{N-1}).$$

Figure 5-2 depicts the function $U_4(q_4, \widehat{h}_1, \widehat{h}_2, \widehat{h}_3)$ for $N = 4$ for eight possible combinations of Alexis’s, Britta’s, and Carol’s decisions $[\widehat{h}_1 \widehat{h}_2 \widehat{h}_3]$. It shows a property of U_n that the updated belief is most dependent on the last public signal \widehat{h}_{n-1} and least on the first public signal \widehat{h}_1 .

An extreme case is when Agent $N - 1$ has not followed precedent. Let us compare two sets of public signals $[1 \ 1 \ 0]$ and $[0 \ 0 \ 1]$: The former leads to a higher belief on $H = 0$ than the latter does even though the latter has more 0. This is because Diana rationally concludes that Carol observed a strong private signal to justify a deviation from precedent.

A reversal of an arbitrarily long precedent sequence may occur because private signals are unbounded; if private signals are bounded like in [17,18], then the influence of precedent can reach a point where agents cannot receive a signal strong enough to justify a decision running counter to precedent (see Section 2.3.2).

■ 5.2 Optimal Initial Belief

The belief update of Agent n discussed in Section 5.1 is the behavior of minimizing the Bayes risk for a given belief q_n . This section discusses what is the effect of incorrect beliefs. In a two-agent case for simplicity, let us find Alexis's and Britta's beliefs that minimize Britta's Bayes risk. For clarification, "Alexis's optimal belief" does not mean the belief optimal for her own sake; it means her belief optimal for Britta.

Britta's Bayes risk R_2 is determined by Alexis's and Britta's error probabilities:

$$\begin{aligned}
R_2 &= c_{10}p_{\widehat{H}_2,H}(1,0) + c_{01}p_{\widehat{H}_2,H}(0,1) \\
&= c_{10}p_H(0) \left[p_{\widehat{H}_2|\widehat{H}_1,H}(1|0,0)p_{\widehat{H}_1|H}(0|0) + p_{\widehat{H}_2|\widehat{H}_1,H}(1|1,0)p_{\widehat{H}_1|H}(1|0) \right] \\
&\quad + c_{01}p_H(1) \left[p_{\widehat{H}_2|\widehat{H}_1,H}(0|0,1)p_{\widehat{H}_1|H}(0|1) + p_{\widehat{H}_2|\widehat{H}_1,H}(0|1,1)p_{\widehat{H}_1|H}(1|1) \right] \\
&= c_{10}p_0 \left[P_{e,2}^{I_0} (1 - P_{e,1}^I) + P_{e,2}^{I_1} P_{e,1}^I \right] + c_{01}(1 - p_0) \left[P_{e,2}^{II_0} P_{e,1}^{II} + P_{e,2}^{II_1} (1 - P_{e,1}^{II}) \right], \quad (5.14)
\end{aligned}$$

where $P_{e,1}^I$ and $P_{e,1}^{II}$ denote Alexis's false alarm and missed detection probabilities and are functions of her belief q_1 . In addition, $P_{e,2}^{I_h}$ and $P_{e,2}^{II_h}$ denote Britta's false alarm and missed detection probabilities when $\widehat{H}_1 = h$ and are functions of her belief q_2 . Britta's error probabilities conditioned on \widehat{H}_1 are not functions of Alexis's belief q_1 even though Britta does social learning with her decision \widehat{H}_1 .

It seems natural that R_2 would be minimum at $q_1 = q_2 = p_0$ because, in this case,

Alexis will make the best decision she can and Britta will not misunderstand her decision. Surprisingly, however, this point is not the minimizer.

Theorem 5.1. *Britta's Bayes risk is not minimum at $q_1 = q_2 = p_0$ in general.*

Proof. Let us find q_1^* where the first derivative of (5.14) with respect to q_1 is zero:

$$\left. \frac{dR_2}{dq_1} \right|_{q_1=q_1^*} = \left[c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0}) \frac{dP_{e,1}^I}{dq_1} + c_{01}(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1}) \frac{dP_{e,1}^{II}}{dq_1} \right]_{q_1=q_1^*} = 0.$$

Let $\lambda_1 = r(q_1)$ denote Alexis's decision threshold. Using

$$\begin{aligned} \frac{dP_{e,1}^I}{dq_1} &= \frac{dP_{e,1}^I}{d\lambda_1} \frac{d\lambda_1}{dq_1} = -f_{Y_1|H}(\lambda_1|0) \frac{d\lambda_1}{dq_1}, \\ \frac{dP_{e,1}^{II}}{dq_1} &= \frac{dP_{e,1}^{II}}{d\lambda_1} \frac{d\lambda_1}{dq_1} = f_{Y_1|H}(\lambda_1|1) \frac{d\lambda_1}{dq_1}, \end{aligned}$$

the first derivative equals zero when

$$\frac{f_{Y_1|H}(\lambda_1|1)}{f_{Y_1|H}(\lambda_1|0)} = \frac{c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})}{c_{01}(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}. \quad (5.15)$$

Since λ_1 is a solution to (5.4), q_1^* satisfies

$$\frac{q_1^*}{1-q_1^*} = \frac{p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})}{(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}. \quad (5.16)$$

Alexis's optimal belief q_1^* is equal to p_0 only when $(P_{e,2}^{I_1} - P_{e,2}^{I_0})/(P_{e,2}^{II_0} - P_{e,2}^{II_1}) = 1$. However, the ratio does not have to be 1. Specifically, in additive Gaussian noise cases, it is not equal to 1 except when $p_0 = c_{01}/(c_{10} + c_{01})$. Therefore, generally speaking,

$$q_1^* \neq p_0.$$

□

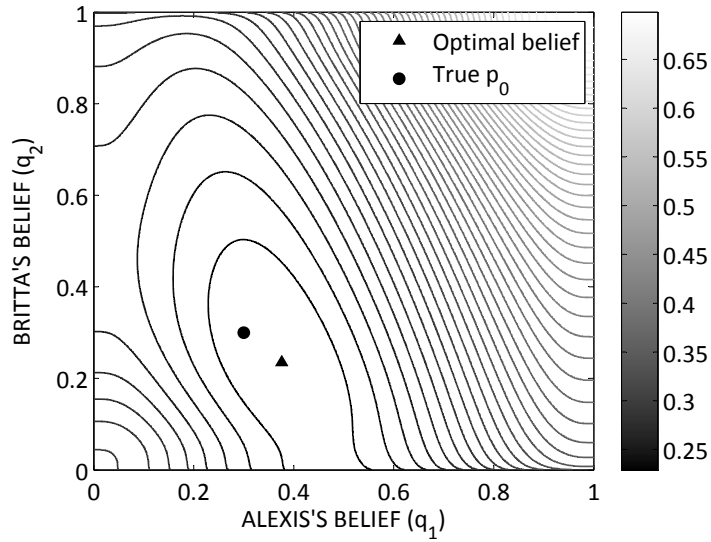


Figure 5-3. Visualization of Britta's Bayes risk for various q_1 and q_2 for $N = 2$, $c_{10} = c_{01} = 1$, $p_0 = 0.3$, and additive Gaussian noise with zero mean and unit variance. Alexis's and Britta's optimal initial beliefs (\blacktriangle) are different from the true prior probability (\bullet).

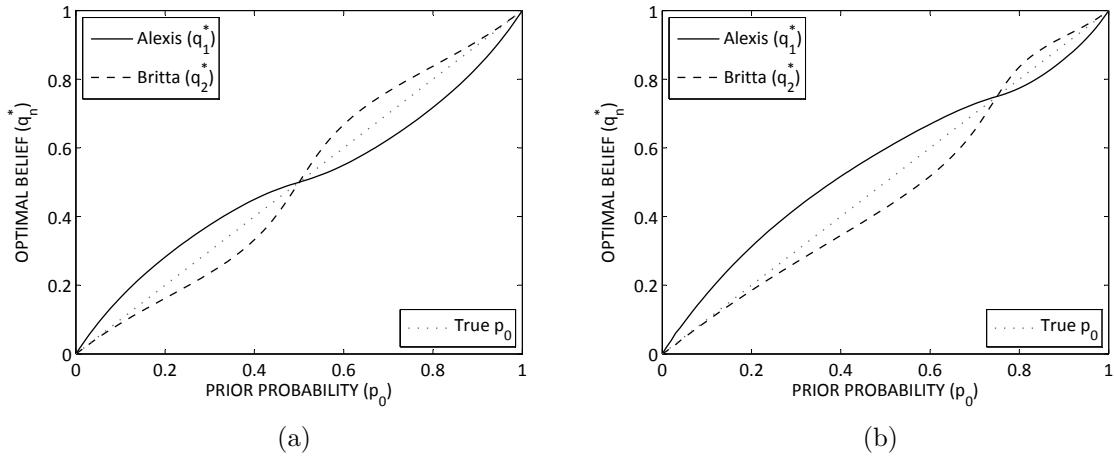


Figure 5-4. The trend of the optimal initial beliefs for varying p_0 for $N = 2$ (Alexis and Britta). (a) $c_{10} = c_{01} = 1$. (b) $c_{10} = 1$, $c_{01} = 3$.

■ 5.2.1 An Open-Minded Advisor

Figure 5-3 depicts Britta's Bayes risk with respect to Alexis's and Britta's beliefs q_1 and q_2 for $\sigma^2 = 1$. Even though the prior probability is 0.3, their optimal beliefs are different from 0.3. Alexis's optimal belief is larger than 0.3 and Britta's belief is smaller than that.

These optimal beliefs follow systematic patterns as shown in Figure 5-4. First,

Alexis should have beliefs larger than p_0 when p_0 is small and beliefs smaller than p_0 when p_0 is large. We call this *open-mindedness* because it is to assign a higher belief to outcomes that are very unlikely. Second, Britta should have a belief smaller than p_0 when p_0 is small and a belief larger than p_0 when p_0 is large. This is necessary to compensate for Alexis's biases. Last, there is a unique fixed point, except for $p_0 = 0$ or $p_0 = 1$, where their optimal beliefs are the same as p_0 . It occurs at $p_0 = c_{01}/(c_{10} + c_{01})$. The following theorems are about the first and the last patterns for $N = 2$.

Theorem 5.2. *Alexis's and Britta's optimal prior beliefs are equal to the true prior probability p_0 if $p_0 = c_{01}/(c_{10} + c_{01})$.*

Proof. The proof is in Appendix 5.A. □

Theorem 5.3. *Let $p_0 \in (0, 1)$ denote the true prior probability and q_1^* Alexis's (i.e., the first agent's) optimal prior belief.*

- *If $p_0 < c_{01}/(c_{10} + c_{01})$, then $p_0 < q_1^* < c_{01}/(c_{10} + c_{01})$.*
- *If $p_0 = c_{01}/(c_{10} + c_{01})$, then $q_1^* = p_0$.*
- *If $p_0 > c_{01}/(c_{10} + c_{01})$, then $c_{01}/(c_{10} + c_{01}) < q_1^* < p_0$.*

Proof. The proof is in Appendix 5.B. □

The patterns also appear in a three-agent case. Figure 5-5 shows agents' beliefs that minimize Carol's Bayes risk. The non-terminal agents (Alexis and Britta) should be open-minded and the last agent (Carol) should be closed-minded. Furthermore, there is a fixed point at the same position as in the two-agent case: $p_0 = c_{01}/(c_{10} + c_{01})$.

■ 5.2.2 An Informative Public Signal

One good thing of having an open-minded advisor is that her decision is more informative than the decision made by a perfect agent. Let us discuss the meaning of being informative. The more private signals an agent has, the better decision she makes. Other than one private signal, however, the agent can only observe others'

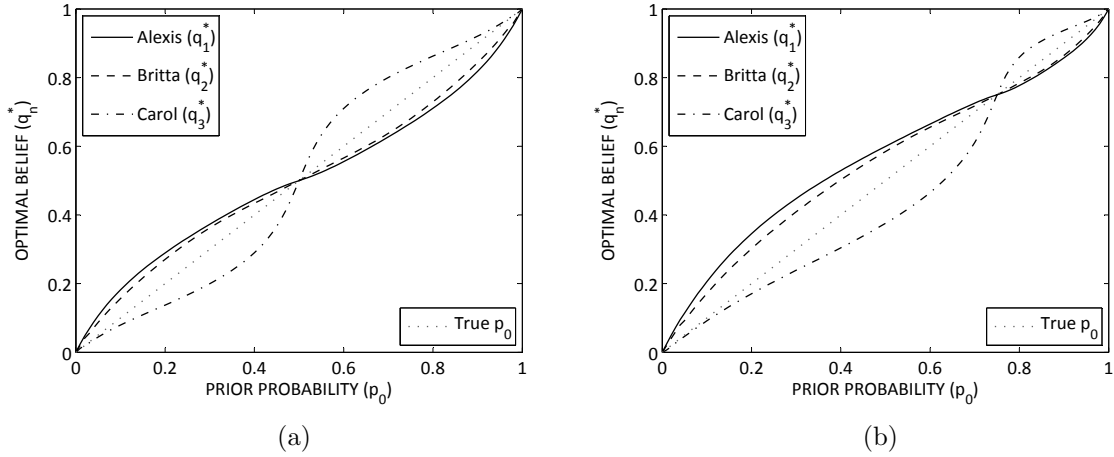


Figure 5-5. The trend of the optimal initial beliefs for varying p_0 for $N = 3$ (Alexis, Britta, and Carol). (a) $c_{10} = c_{01} = 1$. (b) $c_{10} = 1$, $c_{01} = 3$.

decisions, which are 1-bit quantized signals. Hence, being informative means for an agent to make a decision that contains more information about H .

Alexis's open-minded belief increases the mutual information $I(H; \widehat{H}_1)$. Figure 5-6 compares the mutual information in the cases when her belief is equal to q_1^* in Figure 5-4a and when it is equal to p_0 . For comparison, the maximum mutual information $\max_{q_1 \in [0,1]} I(H; \widehat{H}_1)$ is depicted as well. Her optimal belief for Britta yields not maximum but higher mutual information than the true prior p_0 does.

Selfless Agents

All agents considered so far are selfish; it was lucky for a later-acting agent to observe decisions made by incidentally open-minded agents. Let us change the perspective here and think of perfect agents who know p_0 but are selfless. They would want to balance between being right for themselves and being informative for the later-acting agents. A feasible utility function of selfless Agent n is a linear combination of the Bayes risk and the mutual information. It is defined as a function of her decision rule $\widehat{H}_n(\cdot)$:

$$u_n(\widehat{H}_n(\cdot)) = -R_n + \alpha_n I(H; \widehat{H}_n), \quad (5.17)$$

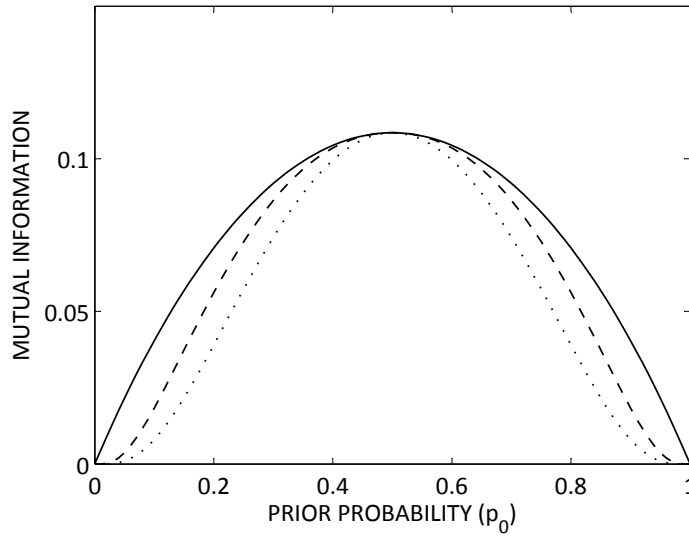


Figure 5-6. Mutual information $I(H; \widehat{H}_1)$ for Alexis's various choices of belief for $N = 2$: From top to bottom, maximum $I(H; \widehat{H}_1)$ (—), her optimal belief for Britta (---), and the belief equal to prior probability (⋯).

where R_n denotes Agent n 's Bayes risk and α_n is a nonnegative parameter that represents the extent of her selflessness. She is selfish when $\alpha_n = 0$ and becomes more selfless as α_n gets bigger.

If Agent n adopts the decision rule that maximizes her utility (5.17), then her behavior will be open-minded. For example, when Alexis has $\alpha_1 = 1$ and uses the decision rule maximizing (5.17), she acts as if she has an open-minded belief that is almost optimal for Britta or Carol. These are depicted in Figure 5-7.

Other-Regarding Preferences

Economic theories has been developed to capture human behaviors that concern others' utilities as well as their own. Such behaviors are called *other-regarding* behaviors. Rabin [25] first proposed a framework of *fairness*: People are willing to sacrifice their own utilities to help those who are being kind and to punish those who are being unkind. He defined two kindness functions in a game of two players i and j :

- A function $f_i(a_i, b_j)$ denotes the kindness of Player i when he chooses a_i and believes that Player j chooses b_j . If Player i gives Player j less than her equitable payoff, $f_i < 0$. Otherwise, $f_i > 0$. If Player i gives Player j her equitable payoff,

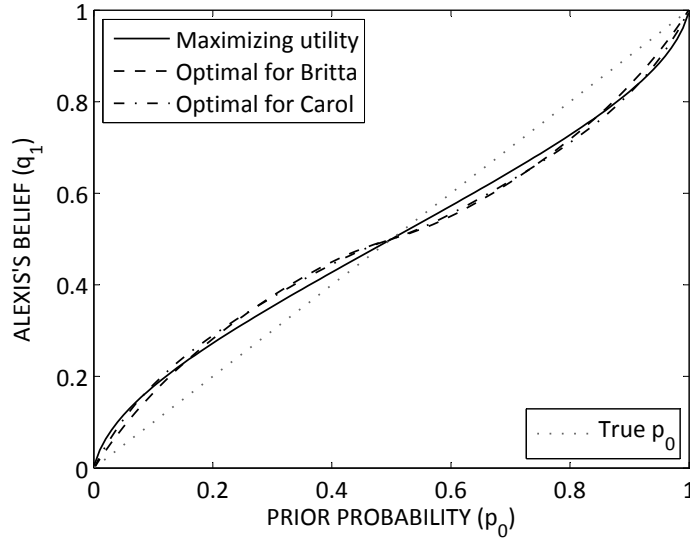


Figure 5-7. Comparison of Alexis’s belief that maximizes her utility (5.17) and her belief that minimizes Britta’s or Carol’s Bayes risk for $c_{10} = c_{01} = 1$.

$f_i = 0$, which indicates that there is no issue of kindness.

- A function $\tilde{f}_i(b_j, c_i)$ denotes Player i ’s belief about how kind Player j is being to him when he believes that Player j believes that his choice is c_i . If Player i believes that Player j is treating him badly, then $\tilde{f}_i < 0$. Otherwise, $\tilde{f}_i > 0$.

The players’ preferences are specified by their utility functions $u_i(a_i, b_j, c_i)$ that include the kindness functions,

$$u_i(a_i, b_j, c_i) = \pi_i(a_i, b_j) + \tilde{f}_i(b_j, c_i) [1 + f_i(a_i, b_j)], \quad (5.18)$$

where the first term $\pi_i(a_i, b_j)$ denotes Player i ’s material utility when he chooses a_i and Player j chooses b_j and the second term is the notion of fairness. If Player i believes that Player j is treating him badly then he will also treat her badly by choosing a_i that makes f_i negative, and vice versa.

Another aspect of other-regarding preferences is *inequity-aversion* [41]. It means that people resist inequitable outcomes; people are willing to give up some utility when they have more utility than others do and, in addition, want to reduce the difference when they have less.

The utility function has the following form:

$$u_i(\pi_1, \pi_2, \dots, \pi_N) = \pi_i - \alpha_i \frac{1}{N-1} \sum_{j \neq i} \max\{\pi_j - \pi_i, 0\} - \beta_i \frac{1}{N-1} \sum_{j \neq i} \max\{\pi_i - \pi_j, 0\}, \quad (5.19)$$

where $\pi_1, \pi_2, \dots, \pi_N$ denote utilities of Players 1, 2, \dots , N , respectively, and α_i and β_i denote Player i 's weights on her advantageous inequality ($\pi_i > \pi_j$) and disadvantageous inequality ($\pi_i < \pi_j$), respectively. In the two-player case, the utility function is simplified to

$$u_i(\pi_i, \pi_j) = \pi_i - \alpha_i \max\{\pi_j - \pi_i, 0\} - \beta_i \max\{\pi_i - \pi_j, 0\}.$$

The utility function (5.18) can capture reciprocal fairness while (5.19) captures distribution of utilities. The combined utility function was proposed in [42]:

$$u_i(\pi_i, \pi_j) = \begin{cases} (1 - (\rho - s\theta))\pi_i + (\rho - s\theta)\pi_j, & \text{if } \pi_i > \pi_j, \\ (1 - (\sigma - s\theta))\pi_i + (\sigma - s\theta)\pi_j, & \text{if } \pi_i < \pi_j, \end{cases} \quad (5.20)$$

where $s = 1$ if Player j has misbehaved and $s = 0$ otherwise. The parameters ρ , σ , and θ capture various aspects of other-regarding preferences. For example, $\sigma \leq \rho \leq 0$ if Player i always prefers to get as much as possible in comparison to Player j , which is called competitive preference. The inequity-aversion corresponds to $\sigma < 0 < \rho < 1$. Furthermore, reciprocity is described by $\theta > 0$. When Player j misbehaves, Player i punishes Player j by reducing both ρ and σ by amount θ .

The selflessness of earlier-acting agents in the sequential detection model is also a kind of other-regarding preferences but different from the reciprocity or distribution preference. First, the detection model does not have the concept of misbehavior. Second, later-acting agents generally achieve lower Bayes risk (i.e., better performance) than earlier-acting agents do because of more affluent public signals. A selfless agent ends up with her own increased Bayes risk to reduce Bayes risk of later-acting agents, which were already smaller than hers. Her behavior increases the inequity between them.

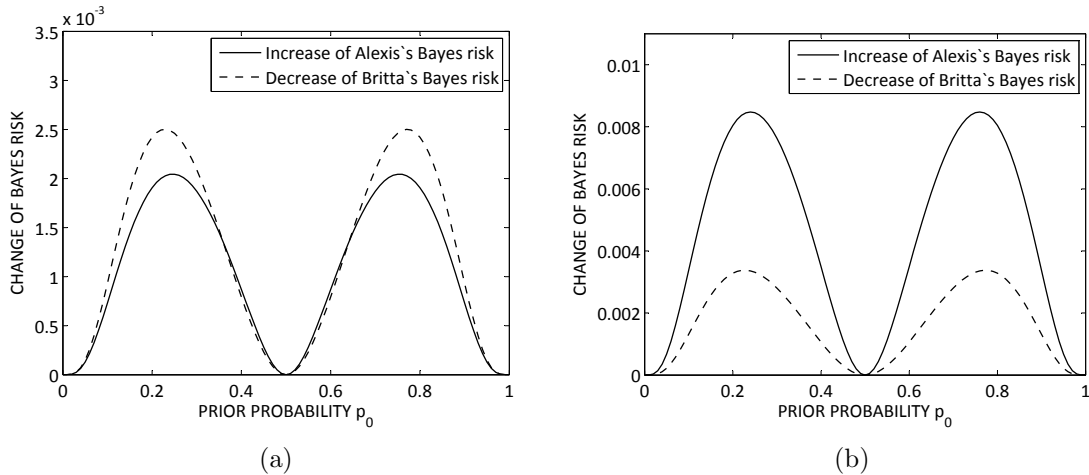


Figure 5-8. Increase of Alexis’s Bayes risk and decrease of Britta’s Bayes risk compared to their Bayes risks when they are perfect agents. (a) Alexis is halfway open-minded, i.e., $q_1 = (p_0 + q_1^*)$. Her loss of performance is smaller than Britta’s gain. (b) Alexis is fully open-minded, i.e., $q_1 = q_1^*$. Her loss of performance is much larger than Britta’s gain. She does not have a reason to be this much open-minded.

Therefore, the selfless utility function (5.17) corresponds to $0 < \sigma < 1$ and $\theta = 0$ in (5.20). In the two-agent case, if Alexis cares both for her own Bayes risk and Britta’s then she will be slightly open-minded for decrease of Britta’s Bayes risk with exchange of small increase of her Bayes risk. However, she would not want to be as open-minded as being optimal for Britta because she has to sacrifice more than Britta’s gain, Figure 5-8

■ 5.3 Conclusion

We have discussed sequential detection performed by a group of agents that make decisions based on individually biased beliefs. Instead of investigating herding on a wrong action, we have assumed unbounded private signals and focused on the agents’ belief update.

The wrong beliefs held by previous agents change the probability with which following agents choose each hypothesis. Contrary to intuition, however, wrong beliefs are not always bad. In fact, the accurate belief—the true prior probability—does not minimize the Bayes risk of the last agent. Specifically, in the case when observations are corrupted by iid additive Gaussian noises, an imperfect agent biased toward

$c_{01}/(c_{10} + c_{01})$ can be more beneficial to subsequent agents than is a perfect agent. In terms of human decision making, where precedent agents are advisors or counselors to the last agent who has the final decisive power, we can say that the best advisors are necessarily open-minded people.

The idea of an open-minded advisor is related to being informative. A public signal is a quantized version of a private signal while also simultaneously reflecting an agent's belief. Alexis's decision will reflect her belief more than her private signal when her belief is very small (close to 0) or very large (close to 1). However, Britta would want a public signal that is most informative of Alexis's private signal. Therefore, she wants her to make her decision with a less extreme belief or an open mind.

While some conclusions of our study depend on having Gaussian likelihoods and may not hold for different types of additive noise, it is more generally true that the optimal beliefs are different from the true prior probability.

■ 5.A Proof of Theorem 5.2

We will show that $\partial R_2/\partial q_1 = 0$ and $\partial R_2/\partial q_2 = 0$ for $q_1 = q_2 = p_0 = c_{01}/(c_{10} + c_{01})$. Then they are Alexis's and Britta's optimal prior beliefs q_1^* and q_2^* .

First, consider $\partial R_2/\partial q_2$ using (5.14):

$$\begin{aligned} \frac{\partial R_2}{\partial q_2} = & -c_{10}p_0 \left[(1 - P_{e,1}^I) f_{Y_2|H}(\lambda_2^0 | 0) \frac{d\lambda_2^0}{dq_2} + P_{e,1}^I f_{Y_2|H}(\lambda_2^1 | 0) \frac{d\lambda_2^1}{dq_2} \right] \\ & + c_{01}(1 - p_0) \left[P_{e,1}^{II} f_{Y_2|H}(\lambda_2^0 | 1) \frac{d\lambda_2^0}{dq_2} + (1 - P_{e,1}^{II}) f_{Y_2|H}(\lambda_2^1 | 1) \frac{d\lambda_2^1}{dq_2} \right]. \end{aligned} \quad (5.21)$$

From (5.5),

$$\frac{q_2^0}{1 - q_2^0} = \frac{q_2}{1 - q_2} \frac{1 - P_{e,1B}^I}{P_{e,1B}^{II}}.$$

Then, for Gaussian likelihoods

$$\begin{aligned}\lambda_2^0 &= r(q_2^0) = \frac{h_1}{2} + \frac{\sigma^2}{h_1} \log \frac{c_{10}q_2^0}{c_{01}(1-q_2^0)} = \left(\frac{h_1}{2} + \frac{\sigma^2}{h_1} \log \frac{c_{10}q_2}{c_{01}(1-q_2)} \right) + \log \frac{1 - P_{e,1B}^I}{P_{e,1B}^{II}} \\ &= r(q_2) + \log \frac{1 - P_{e,1B}^I}{P_{e,1B}^{II}} = \lambda_{1B} + \log \frac{1 - P_{e,1B}^I}{P_{e,1B}^{II}},\end{aligned}$$

and its derivative is given by

$$\begin{aligned}\frac{d\lambda_2^0}{dq_2} &= \frac{d\lambda_{1B}}{dq_2} - \frac{dP_{e,1B}^I}{dq_2} \frac{1}{1 - P_{e,1B}^I} - \frac{dP_{e,1B}^{II}}{dq_2} \frac{1}{P_{e,1B}^{II}} \\ &= \left[1 + \frac{f_{Y_1|H}(\lambda_{1B}|0)}{1 - P_{e,1B}^I} - \frac{f_{Y_1|H}(\lambda_{1B}|1)}{P_{e,1B}^{II}} \right] \frac{d\lambda_{1B}}{dq_2}.\end{aligned}\quad (5.22)$$

Likewise,

$$\lambda_2^1 = \lambda_{1B} + \log \frac{P_{e,1B}^I}{1 - P_{e,1B}^{II}},$$

and its derivative is

$$\frac{d\lambda_2^1}{dq_2} = \left[1 - \frac{f_{Y_1|H}(\lambda_{1B}|0)}{P_{e,1B}^I} + \frac{f_{Y_1|H}(\lambda_{1B}|1)}{1 - P_{e,1B}^{II}} \right] \frac{d\lambda_{1B}}{dq_2}.\quad (5.23)$$

In addition, $q_1 = q_2$ implies that $P_{e,1B}^I = P_{e,1}^I$ and $P_{e,1B}^{II} = P_{e,1}^{II}$, and we can derive the following relations for $q_1 = q_2 = p_0$:

$$\begin{aligned}\frac{f_{Y_2|H}(\lambda_2^0|1)}{f_{Y_2|H}(\lambda_2^0|0)} &= \frac{c_{10}q_2(1 - P_{e,1B}^I)}{c_{01}(1 - q_2)P_{e,1B}^{II}} = \frac{c_{10}p_0(1 - P_{e,1}^I)}{c_{01}(1 - p_0)P_{e,1}^{II}}, \\ \frac{f_{Y_2|H}(\lambda_2^1|1)}{f_{Y_2|H}(\lambda_2^1|0)} &= \frac{c_{10}q_2P_{e,1B}^I}{c_{01}(1 - q_2)(1 - P_{e,1B}^{II})} = \frac{c_{10}p_0P_{e,1}^I}{c_{01}(1 - p_0)(1 - P_{e,1}^{II})}.\end{aligned}\quad (5.24)$$

By substituting (5.22) and (5.23) in (5.21) and using the relations (5.24), we obtain that $\partial R_2/\partial q_2 = 0$ at $q_1 = q_2 = p_0$.

Next, we consider $\partial R_2/\partial q_1$, which is zero at q_1 and q_2 that satisfy (5.16),

$$\frac{q_1}{(1 - q_1)} = \frac{p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})}{(1 - p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}.$$

The condition $q_2 = c_{01}/(c_{10} + c_{01})$ leads to $\lambda_{1B} = h_1/2$ and $P_{e,1B}^I = P_{e,1B}^{II}$. Hence, $\lambda_2^0 - \lambda_{1B} = \lambda_{1B} - \lambda_2^1$, which is equivalent to $\lambda_2^0 = h_1 - \lambda_2^1$. The error probabilities are computed as follows: For $\widehat{H}_1 = 0$,

$$P_{e,2}^{I_0} = p_{\widehat{H}_2|\widehat{H}_1,H}(1|0,0) = \int_{\lambda_2^0}^{\infty} f_{Y_2|H}(y|0) dy,$$

$$P_{e,2}^{II_0} = p_{\widehat{H}_2|\widehat{H}_1,H}(0|0,1) = \int_{-\infty}^{\lambda_2^0} f_{Y_2|H}(y|1) dy,$$

and, for $\widehat{H}_1 = 1$,

$$P_{e,2}^{I_1} = p_{\widehat{H}_2|\widehat{H}_1,H}(1|1,0) = \int_{\lambda_2^1}^{\infty} f_{Y_2|H}(y|0) dy,$$

$$P_{e,2}^{II_1} = p_{\widehat{H}_2|\widehat{H}_1,H}(0|1,1) = \int_{-\infty}^{\lambda_2^1} f_{Y_2|H}(y|1) dy.$$

From the relationship $\lambda_2^0 = h_1 - \lambda_2^1$, $P_{e,2}^{I_0} = P_{e,2}^{II_1}$ and $P_{e,2}^{I_1} = P_{e,2}^{II_0}$ are obtained. Therefore, only $q_2 = p_0$ completes (5.16) and makes $\partial R_2/\partial q_1$ zero.

■ 5.B Proof of Theorem 5.3

The private signals are drawn from the likelihood functions (5.1), but, for the proof of a more general case, let us consider an arbitrary positive number h_1 instead of $H = 1$:

$$f_{Y|H}(y|0) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{y^2}{2\sigma^2}\right), \quad (5.25a)$$

$$f_{Y|H}(y|1) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(y-h_1)^2}{2\sigma^2}\right). \quad (5.25b)$$

Conjecture 5.4. *If $\lambda < h_1/2$, then*

$$\int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2} + \lambda h_1\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2} + \lambda h_1\right) dy$$

$$< \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2} + y h_1\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2} + y h_1\right) dy. \quad (5.26)$$




Figure 5-9. The difference between the left-hand side and the right-hand side of (5.26). (a) $h_1 = 1$. (b) $h_1 = 0.5$.

Figure 5-9 depicts the difference between the left-hand side and the right-hand side of (5.26) for two values of h_1 and supports the conjecture. In the following we assume the conjecture to be true.

Lemma 5.5. *If $\lambda < h_1/2$, then*

$$\begin{aligned}
 & \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{\lambda h_1}{\sigma^2}\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{\lambda h_1}{\sigma^2}\right) dy \\
 & < \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{y h_1}{\sigma^2}\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{y h_1}{\sigma^2}\right) dy.
 \end{aligned} \tag{5.27}$$

Proof. Substituting $y' = y/\sigma$, $\lambda' = \lambda/\sigma$, and $h'_1 = h_1/\sigma$, we obtain

$$\begin{aligned}
 & \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{\lambda h_1}{\sigma^2}\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{\lambda h_1}{\sigma^2}\right) dy \\
 & = \sigma^2 \int_{-\infty}^{\lambda'} \exp\left(-\frac{y'^2}{2} + \lambda' h'_1\right) dy' \int_{\lambda'}^{\infty} \exp\left(-\frac{y'^2}{2} + \lambda' h'_1\right) dy'
 \end{aligned}$$

and

$$\begin{aligned}
 & \int_{-\infty}^{\lambda} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{y h_1}{\sigma^2}\right) dy \int_{\lambda}^{\infty} \exp\left(-\frac{y^2}{2\sigma^2} + \frac{y h_1}{\sigma^2}\right) dy \\
 & = \sigma^2 \int_{-\infty}^{\lambda'} \exp\left(-\frac{y'^2}{2} + y' h'_1\right) dy' \int_{\lambda'}^{\infty} \exp\left(-\frac{y'^2}{2} + y' h'_1\right) dy'.
 \end{aligned}$$

Since $\lambda < h_1/2$ implies $\lambda' < h_1'/2$, (5.27) follows from Conjecture 5.4. \square

Here the proof of Theorem 5.3 begins: First, the case when $p_0 = c_{01}/(c_{10} + c_{01})$ is proven in Appendix 5.A.

Next, suppose that $p_0 < c_{01}/(c_{10} + c_{01})$. Let λ in (5.27) denote a decision threshold according to q_2^* . Obviously, optimal prior beliefs q_1^* and q_2^* should be strictly decreasing as p_0 decreases like in Figure 5-2. Hence, Theorem 5.2, which states that $q_1^* = q_2^* = c_{01}/(c_{10} + c_{01})$ if $p_0 = c_{01}/(c_{10} + c_{01})$, implies that

$$q_1^* < \frac{c_{01}}{c_{10} + c_{01}} \quad \text{and} \quad q_2^* < \frac{c_{01}}{c_{10} + c_{01}} \quad (5.28)$$

if $p_0 < c_{01}/(c_{10} + c_{01})$. Then we get $\lambda < h_1/2$ and can use (5.27).

Multiplying each integrand in (5.27) by the constant $\frac{1}{2\pi\sigma^2} \exp(-\frac{\lambda^2}{2\sigma^2} - \frac{h_1^2}{2\sigma^2})$, we get

$$\begin{aligned} & \int_{-\infty}^{\lambda} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(\lambda - h_1)^2}{2\sigma^2}\right) \exp\left(-\frac{y^2}{2\sigma^2}\right) dy \\ & \times \int_{\lambda}^{\infty} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(\lambda - h_1)^2}{2\sigma^2}\right) \exp\left(-\frac{y^2}{2\sigma^2}\right) dy \\ & < \int_{-\infty}^{\lambda} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\lambda^2}{2\sigma^2}\right) \exp\left(-\frac{(y - h_1)^2}{2\sigma^2}\right) dy \\ & \times \int_{\lambda}^{\infty} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{\lambda^2}{2\sigma^2}\right) \exp\left(-\frac{(y - h_1)^2}{2\sigma^2}\right) dy. \end{aligned} \quad (5.29)$$

According to (5.25), the exponential functions in (5.29) are likelihood functions of Y_1 , so we have

$$\begin{aligned} & f_{Y_1|H}^2(\lambda|1) \int_{-\infty}^{\lambda} f_{Y_1|H}(y|0) dy \int_{\lambda}^{\infty} f_{Y_1|H}(y|0) dy \\ & < f_{Y_1|H}^2(\lambda|0) \int_{-\infty}^{\lambda} f_{Y_1|H}(y|1) dy \int_{\lambda}^{\infty} f_{Y_1|H}(y|1) dy. \end{aligned} \quad (5.30)$$

Since $\lambda = r(q_2^*)$, we obtain

$$\frac{P_{e,1B}^I (1 - P_{e,1B}^I)}{f_{Y_1|H}^2(\lambda|0)} < \frac{P_{e,1B}^{II} (1 - P_{e,1B}^{II})}{f_{Y_1|H}^2(\lambda|1)}, \quad (5.31)$$

and (5.4) transforms (5.31) to

$$c_{10}^2 q_2^{*2} P_{e,1B}^I (1 - P_{e,1B}^I) < c_{01}^2 (1 - q_2^*)^2 P_{e,1B}^{II} (1 - P_{e,1B}^{II}).$$

Rearranging terms gives us

$$\left(\frac{c_{10} q_2^*}{c_{01} (1 - q_2^*)} \frac{1 - P_{e,1B}^I}{P_{e,1B}^{II}} \right)^{-1} > \frac{c_{10} q_2^*}{c_{01} (1 - q_2^*)} \frac{P_{e,1B}^I}{1 - P_{e,1B}^{II}}. \quad (5.32)$$

The terms in the left-hand and the right-hand sides are the same as the updated beliefs $q_2^0/(1 - q_2^0)$ and $q_2^1/(1 - q_2^1)$, respectively, so this simplifies to

$$\left(\frac{c_{10} q_2^0}{c_{01} (1 - q_2^0)} \right)^{-1} > \frac{c_{10} q_2^1}{c_{01} (1 - q_2^1)}. \quad (5.33)$$

Let us discuss the meaning of the inequality (5.33). In Figure 5-10, the convex curve depicts a flipped version of the ROC curve. When the prior belief is q , the error probabilities $(P_{e,2}^I, P_{e,2}^{II})$ are determined as the point of tangency, where the curve meets a line with slope $-c_{10}q/c_{01}(1 - q)$. Two solid lines in Figure 5-10 depict the lines for Britta's updated beliefs after observing $\widehat{H}_1 = 0$ and $\widehat{H}_1 = 1$, respectively denoted by q_2^0 and q_2^1 .

The inequality (5.33) restricts the range of error probabilities in which $(P_{e,2}^{I_0}, P_{e,2}^{II_0})$ can exist on the basis of $(P_{e,2}^{I_1}, P_{e,2}^{II_1})$; the point $B^0 (P_{e,2}^{I_0}, P_{e,2}^{II_0})$, a black dot, always exists on the right side of the point $\bar{B}^1 (P_{e,2}^{II_1}, P_{e,2}^{I_1})$, a gray diamond. Furthermore, the point B^0 cannot exist on the right side of the point $B^1 (P_{e,2}^{I_1}, P_{e,2}^{II_1})$, a black diamond, because obviously $q_2^0 > q_2^1$. Therefore, the point B^0 always exists on the curve between the points \bar{B}^1 and B^1 .

Now we draw a black dotted line that connects the points B^0 and B^1 and a gray dashed line that connects the points \bar{B}^1 and B^1 . From the restriction for the point B^0 , the slope of the former is always greater than that of the latter:

$$\frac{P_{e,2}^{II_1} - P_{e,2}^{II_0}}{P_{e,2}^{I_1} - P_{e,2}^{I_0}} > -1. \quad (5.34)$$




Figure 5-10. The point $B^0 (P_{e,2}^{I_0}, P_{e,2}^{II_0})$ always exists between the points $\bar{B}^1 (P_{e,2}^{II_1}, P_{e,2}^{I_1})$ and $B^1 (P_{e,2}^{I_1}, P_{e,2}^{II_1})$.

We have obtained the optimality condition (5.16) for Alexis's prior belief in Section 5.2. We can rewrite it as

$$q_1^* = \frac{p_0}{p_0 + (1 - p_0) \frac{P_{e,2}^{II_0} - P_{e,2}^{II_1}}{P_{e,2}^{I_1} - P_{e,2}^{I_0}}}. \quad (5.35)$$

Finally, we can conclude that $q_1^* > p_0$ because of (5.34).

In addition, Alexis's optimal belief q_1^* is upper-bounded by $c_{01}/(c_{10} + c_{01})$ because q_1^* is strictly decreasing in p_0 and $q_1^* = c_{01}/(c_{10} + c_{01})$ when $p_0 = c_{01}/(c_{10} + c_{01})$ by Theorem 5.2. Combining these two bounds, we have the inequality

$$p_0 < q_1^* < \frac{c_{01}}{c_{10} + c_{01}}, \quad (5.36)$$

as desired.

The statement that $c_{01}/(c_{10} + c_{01}) < q_1^* < p_0$ if $p_0 > c_{01}/(c_{10} + c_{01})$ can be proven similarly.

Distributed Detection with Social Learning and Symmetric Fusion

Social learning among individualistically decision-making agents is a rational behavior to decrease Bayes risk. Some imperfect agents having incorrect beliefs can decrease Bayes risk even below the Bayes risk of perfect agents by doing social learning as discussed in Chapter 5.

On the other hand, social learning may be futile within a team of agents who make a decision together. This chapter investigates the effect of social learning within a distributed detection model discussed in Chapter 4. The agents make local decisions in a predetermined order and later-acting agents can observe decisions made by earlier-acting agents. Once all agents make local decisions, the global decision is determined according to an L -out-of- N fusion rule.

Social learning turns out not to be beneficial when the agents observe conditionally iid private signals. In contrast, when they observe private signals not conditionally iid, social learning can be beneficial. Furthermore, the order in which they make local decisions can affect the performance of the team similarly to the work [43], which showed that the order in which heterogeneous agents speak matters to their individual performances. The importance of ordering is also pointed out in [44].

Although, for convenience, private signals are assumed to be corrupted by additive Gaussian noises, most discussions in this chapter do not require the assumption. The only condition required in the discussions is that the likelihood ratios of the private signal under $H = 0$ and $H = 1$ are strictly monotonic, which is a rather natural




Figure 6-1. The distributed detection and decision fusion model considered in this chapter. Local decisions are sent not only to the fusion center but also to following agents as public signals.

requirement.

Section 6.1 provides results for agents observing conditionally iid private signals and all public signals. Section 6.2 adds a constraint that each agent can only observe a subset of the public signals. However, the results will be the same as in Section 6.1. Section 6.3 generalizes the model to agents with differing private signal likelihoods and provides examples of helpful social learning. Section 6.4 explains our experiments to test whether people rationally ignore public signals in distributed detection scenarios discussed in Section 6.1. Section 6.5 summarizes our results and discusses limitations and extensions of our model.

■ 6.1 Conditionally IID Private Signals and Complete Public Signals

The agents, Alexis, Britta, Carol, \dots , Norah, make a decision together in the alphabetical order by voting and a symmetric fusion rule, Figure 6-1. Suppose that they know the prior probability p_0 . Agent n observes a conditionally iid private signal $Y_n = H + W_n$, where $W_n \sim \mathcal{N}(0, \sigma^2)$ and W_n is independent of W_m for any $m \neq n$. She makes a local decision $\hat{H}_n \in \{0, 1\}$, which is sent to a fusion center and also observed by the other agents. The global decision is determined by the fusion center according

to an L -out-of- N fusion rule:

$$\widehat{H} = \begin{cases} 1, & \text{if } \sum_{n=1}^N \widehat{H}_n \geq L; \\ 0, & \text{otherwise,} \end{cases} \quad (6.1)$$

Since we are assuming monotonically increasing likelihood ratios, we can express the LRT in a compact form with a decision threshold λ_n :

$$y_n \underset{\widehat{H}_n(y_n)=0}{\overset{\widehat{H}_n(y_n)=1}{\gtrless}} \lambda_n. \quad (6.2)$$

The global decision is of our interest. The goodness criterion for the decision making is the Bayes risk according to the global decision:

$$R = c_{10}p_0\mathbb{P}\{\widehat{H} = 1 | H = 0\} + c_{01}(1 - p_0)\mathbb{P}\{\widehat{H} = 0 | H = 1\}. \quad (6.3)$$

The agents are a team in the sense of Radner [45]: they share the decision-making cost and the same goal, which is to minimize (6.3).

In this chapter, the optimal performance in the case when agents can observe public signals—called *public voting*—is compared to the case when they cannot—called *secret voting*. It will be proven that the optimal decision thresholds are the same in both cases. The same optimal thresholds imply the same performances because the criterion functions are the same in the public and the secret voting.

■ 6.1.1 Two-Agent Case

Suppose that Alexis and Britta make a decision together and their fusion rule is the 1-out-of-2 (OR) rule.

In secret voting, Alexis and Britta simultaneously make local decisions. They use one decision threshold each. In public voting, Alexis first makes a decision and then Britta makes a decision upon observing Alexis's decision. Alexis uses one decision threshold because she does not observe a public signal anyway.

Britta can have two decision thresholds in public voting according to Alexis's

decision \widehat{H}_1 . However, her decision rule for $\widehat{H}_1 = 1$ is irrelevant to the global decision because of the fusion rule: the global decision is 1 whenever $\widehat{H}_1 = 1$. She makes a pivotal vote only when $\widehat{H}_1 = 0$.

Likewise, if the fusion rule is the 2-out-of-2 (AND) rule, Britta's decision is relevant only when $\widehat{H}_1 = 0$. The social learning degenerates Britta's decision making; she has one relevant decision threshold as in secret voting. Therefore, social learning has no merit in the two-agent case.

This is an intuition of the statement that public voting and secret voting yield the same performance for $N = 2$. Theorem 6.1 provides a more technical proof. Throughout this chapter, discriminable notations are used to denote agents' decision thresholds in public voting and in secret voting: λ in secret voting and ρ in public voting.

Theorem 6.1. *The existence of public signals does not affect the optimal local decision rules for $N = 2$ under either of the two L -out-of- N fusion rules.*

Proof. Let us compare team Bayes risks of secret voting and public voting under the OR rule.

In the secret voting scenario, the Bayes risk is given by

$$R_s = c_{10}p_0 \left(P_{e,1}^I + (1 - P_{e,1}^I) P_{e,2}^I \right) + c_{01}(1 - p_0) P_{e,1}^{II} P_{e,2}^{II}, \quad (6.4)$$

where Alexis's decision threshold λ_1 determines local error probabilities $P_{e,1}^I$ and $P_{e,1}^{II}$, and Britta's decision threshold λ_2 determines local error probabilities $P_{e,2}^I$ and $P_{e,2}^{II}$. Their optimal decision thresholds λ_1^* and λ_2^* minimize (6.4).

In the public voting scenario, the Bayes risk has the same form

$$R_p = c_{10}p_0 \left(P_{e,1}^I + (1 - P_{e,1}^I) P_{e,2}^{I_0} \right) + c_{01}(1 - p_0) P_{e,1}^{II} P_{e,2}^{II_0}, \quad (6.5)$$

except that Britta's error probabilities, $P_{e,2}^{I_0}$ and $P_{e,2}^{II_0}$, are controlled by her decision threshold ρ_2^0 for $\widehat{H}_1 = 0$. Britta's decision threshold when $\widehat{H}_1 = 1$, ρ_2^1 , is irrelevant; thus, we can assume that $\rho_2^1 = \rho_2^0$ without loss of optimality.

Expressions (6.4) and (6.5) are very similar; the only difference is the replacement of $(P_{e,2}^I, P_{e,2}^{II})$ in (6.4) by $(P_{e,2}^{I_0}, P_{e,2}^{II_0})$ in (6.5). Now note that the set of achievable values for $(P_{e,2}^I, P_{e,2}^{II})$ and $(P_{e,2}^{I_0}, P_{e,2}^{II_0})$ are identical; they are achieved by varying Britta's decision threshold (λ_2 or ρ_2^0) in precisely the same local decision-making problem. Therefore, minimizing R_s and R_p results in equal Bayes risks, and these are achieved with decision thresholds satisfying the following:

$$\rho_1^* = \lambda_1^*, \quad \rho_2^* = \lambda_2^*. \quad (6.6)$$

Since $\lambda_1^* = \lambda_2^*$, we also have $\rho_1^* = \rho_2^*$. Therefore, Alexis and Britta should not change their decision thresholds depending on whether or not the voting is public.

The proof for the AND fusion rule is similar. The Bayes risks of secret voting and public voting have identical forms

$$\begin{aligned} R_s &= c_{10}p_0P_{e,1}^I P_{e,2}^I + c_{01}(1-p_0) \left(P_{e,1}^I + (1 - P_{e,1}^{II}) P_{e,2}^{II} \right), \\ R_p &= c_{10}p_0P_{e,1}^I P_{e,2}^{I_1} + c_{01}(1-p_0) \left(P_{e,1}^{II} + (1 - P_{e,1}^{II}) P_{e,2}^{II_1} \right). \end{aligned}$$

Again we find that the achievable set of local Type I and Type II error probabilities are identical under secret and public voting, so the minima and optimum decision thresholds are unaffected by the public signal. \square

Theorem 6.1 states that the agents need to use the same decision thresholds whether Britta can or cannot observe Alexis's decision. Therefore, Britta does not need to do social learning when $N = 2$.

■ 6.1.2 N -Agent Case

The comparison is obvious in the two-agent case because the Bayes risk formula is identical in both secret and public voting. However, the Bayes risk formula of secret voting is not the same as that of public voting for $N \geq 3$. For $N = 3$ and the majority fusion rule, for example, Britta's decision is relevant whether $\widehat{H}_1 = 0$ or $\widehat{H}_1 = 1$. Thus, she has two concrete degrees of freedom with respect to choosing decision thresholds

in public voting while she has only one degree of freedom in secret voting.

Theorem 6.2. *Suppose that sharing local decisions does not change Alexis's decision threshold (i.e., $\rho_1^* = \lambda_1^*$). If the optimal decision-making rules are the same in public voting and in secret voting for a specific N and any K -out-of- N fusion rule, then it is also true for a team of $N + 1$ agents and any L -out-of- $(N+1)$ fusion rule.*

Proof. First, we consider the secret voting scenario with $N + 1$ agents. Since Agent n 's decision is critical only if the other N local decisions are $L - 1$ ones and $N - L + 1$ zeros, the optimal decision threshold λ_n^* is the solution to¹

$$\begin{aligned} \frac{f_{Y_n|H}(\lambda_n | 1)}{f_{Y_n|H}(\lambda_n | 0)} &= \frac{c_{10}p_0 \binom{N}{L-1} (P_e^I)^{L-1} (1 - P_e^I)^{N-L+1}}{c_{01}(1-p_0) \binom{N}{N-L+1} (P_e^{II})^{N-L+1} (1 - P_e^{II})^{L-1}} \\ &= \frac{c_{10}p_0 (P_e^I)^{L-1} (1 - P_e^I)^{N-L+1}}{c_{01}(1-p_0) (P_e^{II})^{N-L+1} (1 - P_e^{II})^{L-1}}, \end{aligned} \quad (6.7)$$

where we use that $P_{e,1}^I = P_{e,2}^I = \dots = P_e^I$ and $P_{e,1}^{II} = P_{e,2}^{II} = \dots = P_e^{II}$ because the optimal decision thresholds of all agents are identical in secret voting.

Next, in the public voting scenario, we can classify error cases depending on Alexis's detection result, e.g., when the true state is 0 and Alexis's decision is correct ($\widehat{H}_1 = 0$), a false alarm occurs if at least L out of the remaining N agents vote for 1. The Bayes risk is given by

$$\begin{aligned} R_p &= c_{10}p_0 (1 - P_{e,1}^I) \mathbb{P} \left\{ \sum_{n=2}^{N+1} \widehat{H}_n \geq L \mid \widehat{H}_1 = H = 0 \right\} \\ &\quad + c_{10}p_0 P_{e,1}^I \mathbb{P} \left\{ \sum_{n=2}^{N+1} \widehat{H}_n \geq L - 1 \mid \widehat{H}_1 = 1, H = 0 \right\} \\ &\quad + c_{01}(1-p_0) P_{e,1}^{II} \mathbb{P} \left\{ \sum_{n=2}^{N+1} \widehat{H}_n \leq L - 1 \mid \widehat{H}_1 = 0, H = 1 \right\} \\ &\quad + c_{01}(1-p_0) (1 - P_{e,1}^{II}) \mathbb{P} \left\{ \sum_{n=2}^{N+1} \widehat{H}_n \leq L - 2 \mid \widehat{H}_1 = H = 1 \right\} \\ &\triangleq R_0 (p_0(1 - P_{e,1}^I) + (1 - p_0)P_{e,1}^{II}) + R_1 (p_0P_{e,1}^I + (1 - p_0)(1 - P_{e,1}^{II})), \end{aligned} \quad (6.8)$$

¹Please see Section 2.2.2 for a description of how (6.7) is derived.

where R_0 and R_1 are specified in (6.10) and (6.11) and we define

$$\begin{aligned} q^0 &\triangleq \frac{p_0(1 - P_{e,1}^I)}{p_0(1 - P_{e,1}^I) + (1 - p_0)P_{e,1}^{II}} = \mathbb{P}\{H = 0 \mid \widehat{H}_1 = 0\}, \\ q^1 &\triangleq \frac{p_0P_{e,1}^I}{p_0P_{e,1}^I + (1 - p_0)(1 - P_{e,1}^{II})} = \mathbb{P}\{H = 0 \mid \widehat{H}_1 = 1\}. \end{aligned} \quad (6.9)$$

When Agents 2, 3, \dots , $N + 1$ observe $\widehat{H}_1 = 0$, their optimal decision strategy is to minimize the term R_0 from (6.8):

$$\begin{aligned} R_0 &= c_{10}q^0\mathbb{P}\left\{\sum_{n=2}^{N+1}\widehat{H}_n^0 \geq L \mid H = 0\right\} \\ &\quad + c_{01}(1 - q^0)\mathbb{P}\left\{\sum_{n=2}^{N+1}\widehat{H}_n^0 \leq L - 1 \mid H = 1\right\}, \end{aligned} \quad (6.10)$$

where the condition $\widehat{H}_1 = 0$ is embedded in the superscript of \widehat{H}_n^0 . Please note that R_0 is the same as the Bayes risk of N agents when the prior probability is q^0 and the fusion is done by the L -out-of- N rule. It implies that the optimal decision thresholds of Agents 2, 3, \dots , $N + 1$ are the same as those of N agents with prior probability q^0 and the L -out-of- N fusion rule.

Likewise, when Agents 2, 3, \dots , $N + 1$ observe $\widehat{H}_1 = 1$, their optimal decision strategy is to minimize the term R_1 from (6.8):

$$\begin{aligned} R_1 &= c_{10}q^1\mathbb{P}\left\{\sum_{n=2}^{N+1}\widehat{H}_n^1 \geq L - 1 \mid H = 0\right\} \\ &\quad + c_{01}(1 - q^1)\mathbb{P}\left\{\sum_{n=2}^{N+1}\widehat{H}_n^1 \leq L - 2 \mid H = 1\right\}. \end{aligned} \quad (6.11)$$

Their optimal decision thresholds are the same as those of N agents with prior probability q^1 and the $(L-1)$ -out-of- N fusion rule. Figure 6-2 depicts the evolution of the problem corresponding to Alexis's decision \widehat{H}_1 .

Let us find the optimal thresholds $\rho_2^{0*}, \rho_3^{0*}, \dots, \rho_{N+1}^{0*}$ in Problem B0 in Figure 6-2. In fact, Problem B0 is also a public voting scenario; agents observe $\widehat{H}_2, \widehat{H}_3$, and so on. However, because of the assumption that the existence of the public signals does not affect optimal decision thresholds of a team of N agents for any K -out-of- N fusion

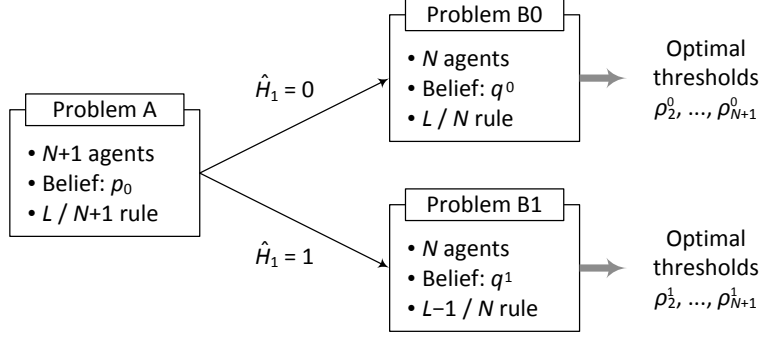


Figure 6-2. An $(N+1)$ -agent problem is divided into two N -agent problems depending on Alexis's decision \hat{H}_1 .

rule, we can find the optimal thresholds as if the agents do secret voting. Since an agent's decision is critical only if the other $N - 1$ local decisions consist of $L - 1$ ones and $N - L$ zeros, the optimal decision threshold ρ^{0*} is the solution to

$$\begin{aligned} \frac{f_{Y|H}(\rho^0 | 1)}{f_{Y|H}(\rho^0 | 0)} &= \frac{c_{10}q^0 \binom{N-1}{L-1} (P_e^{I_0})^{L-1} (1-P_e^{I_0})^{N-L}}{c_{01}(1-q^0) \binom{N-1}{N-L} (P_e^{II_0})^{N-L} (1-P_e^{II_0})^{L-1}} \\ &= \frac{c_{10}p_0 (1-P_e^I) (P_e^{I_0})^{L-1} (1-P_e^{I_0})^{N-L}}{c_{01}(1-p_0) P_e^{II} (P_e^{II_0})^{N-L} (1-P_e^{II_0})^{L-1}}, \end{aligned} \quad (6.12)$$

where q^0 is replaced by (6.9). Due to the assumption that $\rho_1^* = \lambda_1^*$, P_e^I and P_e^{II} in (6.12) are the same as P_e^I and P_e^{II} in (6.7).

Comparing (6.12) to (6.7), we can find that they have the same solutions, i.e., $\rho_i^{0*} = \lambda_i^*$. Therefore, the agents should not change their decision thresholds after observing $\hat{H}_1 = 0$.

We can also find the optimal thresholds $\rho_2^*, \dots, \rho_{N+1}^*$ in Problem B1 in Figure 6-2 by looking at the N -agent problem without public signals:

$$\begin{aligned} \frac{f_{Y|H}(\rho^1 | 1)}{f_{Y|H}(\rho^1 | 0)} &= \frac{c_{10}q^1 \binom{N-1}{L-2} (P_e^{I_1})^{L-2} (1-P_e^{I_1})^{N-L+1}}{c_{01}(1-q^1) \binom{N-1}{N-L+1} (P_e^{II_1})^{N-L+1} (1-P_e^{II_1})^{L-2}} \\ &= \frac{c_{10}p_0 P_e^I (P_e^{I_1})^{L-2} (1-P_e^{I_1})^{N-L+1}}{c_{01}(1-p_0) (1-P_e^{II}) (P_e^{II_1})^{N-L+1} (1-P_e^{II_1})^{L-2}}. \end{aligned} \quad (6.13)$$

Again, due to the assumption that $\rho_1^* = \lambda_1^*$, P_e^I and P_e^{II} in (6.13) are the same as P_e^I

and P_e^{II} in (6.7). We reach the same conclusion that the two equations have the same solutions, i.e., $\rho_i^{1*} = \lambda_i^*$, by comparing (6.13) to (6.7). Thus, the agents should not change their decision thresholds after observing $\widehat{H}_1 = 1$.

Consequently, for a team of $N + 1$ agents and any L -out-of- $(N+1)$ rule, their optimal decision thresholds are the same whether they observe previous decisions or not. \square

Corollary 6.3. *Suppose that sharing local decisions does not change Alexis’s decision rule (i.e., $\rho_1^* = \lambda_1^*$). For any N and L -out-of- N fusion rule, the existence of the public signals does not affect optimal decision thresholds of a team of N agents.*

Proof. Use mathematical induction with Theorems 6.1 and 6.2. \square

This result requires the assumption that Alexis uses the same decision rule in both secret and public voting. This assumption is trivially true for $N = 1$. It is also true for $N = 2$ by the proof of Theorem 6.1. In addition, our numerical experiments for $N \leq 9$ confirmed that it is true. Thus, this assumption seems heuristically true. In particular, we fail to see how Alexis would choose between increasing or decreasing her decision threshold based on the existence of public signals.

Note that the updated belief (6.9) upon observing Alexis’s decision \widehat{H}_1 in distributed detection model is the same as that (5.5) in the individualistic sequential detection model in Chapter 5. It implies that the social learning within a team biases later-acting agents in favor of the dominant choice of earlier-acting agents as much as social learning between individual agents does.

The difference is that *fusion rule evolution* arises from social learning within the team along with the *belief update*. While the belief update is common in the social learning literature—in fact, central to it—the evolution of the fusion rule hardly appears in the literature.

The fusion rule evolution matters because the change of the fusion rule affects computation of the optimal decision threshold. Since all agents can observe all public signals in a public voting scenario, each agent can keep a running tally of the numbers of 0 and 1 votes throughout the voting. If most of earlier-acting agents chose 1, it

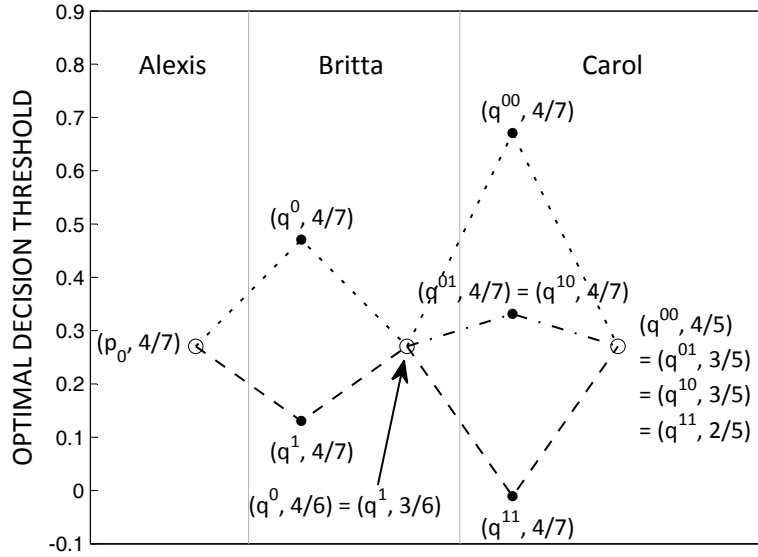


Figure 6-3. Illustration of the separate effects of the belief update and the fusion rule evolution. The diagram depicts decision thresholds for the first three agents in an example with $p_0 = 0.25$ and the 4-out-of-7 fusion rule. Agents observe $H \in \{0, 1\}$ corrupted by iid Gaussian additive noises with zero mean and unit variance, and $c_{10} = c_{01} = 1$. The $(p, L/N)$ label on a marker (\bullet or \circ) indicates that its height represents the optimal decision threshold for prior probability p and the L -out-of- N fusion rule. Alexis has an initial decision threshold depending only on the prior probability p_0 and the 4-out-of-7 fusion rule (leftmost \circ). If Britta considers belief updates only, the optimal decision threshold is changed from Alexis's decision threshold to a new value that depends on \widehat{H}_1 (two leftmost \bullet 's). However, after adopting the fusion rule evolution as well, Britta's optimal decision threshold returns to equal Alexis's decision threshold (center \circ). Similarly, if Carol considers belief updates only, the four values for $(\widehat{H}_1, \widehat{H}_2)$ lead to three distinct decision thresholds (three rightmost \bullet 's). After accounting for the fusion rule evolution, Carol's optimal decision threshold returns to equal Alexis's decision threshold (rightmost \circ).

would imply that only a few more 1 votes are sufficient to close the voting with the global decision as 1. Hence, in order to vote 1, an agent should need a strong private signal in support of 1 enough to take the risk of making later-acting agents' votes irrelevant. Unlike the belief update, the fusion rule evolution discourages later-acting agents to follow the dominant choice of earlier-acting agents.

By Corollary 6.3, the effects of the belief update and fusion rule evolution exactly cancel out. A numerical example is detailed in Figure 6-3. Reading that figure from left to right, for each agent after Alexis, the belief update is done first and then the fusion rule evolution brings the optimal decision threshold back exactly to the optimal decision threshold of Alexis.

In conclusion, social learning within a team of agents causes two effects: Belief

update plays a role of a positive feedback while fusion rule evolution does a negative feedback. By proving that the optimal thresholds of the last N agents are the same in the $(N+1)$ -agent cases when $\widehat{H}_1 = 0$, when $\widehat{H}_1 = 1$, and when they do not know \widehat{H}_1 , it is proved that the effects of the former and the latter are exactly canceled out.

■ 6.1.3 N -Imperfect-Agent Case

Corollary 6.3 is the result in the case when all agents are perfect in the sense that they know the prior probability p_0 . However, it also implies that the Bayes risk of a team of imperfect agents must be higher than that of the team of perfect agent.

It is derived in Section 2.2 that the optimal local decision rule is given by (2.9) for the distributed detection problem with a symmetric fusion rule and without social learning. This is the optimal rule whether the agents know p_0 or not.

Corollary 6.3 asserts that the optimal local decision rule is the same as (2.9) even when the agents are allowed to do social learning. Again, whether the agents know p_0 or not is irrelevant to the optimal threshold. If the agents are imperfect, their rational decision rule will be different from the optimal decision rule. Therefore, the imperfect agents will be always outperformed by perfect agents of the same number if the agents observe conditionally iid private signals.

■ 6.2 Conditionally IID Private Signals and Incomplete Public Signals

The previous section assumes that each and every agent observes decisions of all the other agents. This section considers more restricted agents who can only observe an arbitrary subset of the previously made decisions. For example, each agent may observe the public signal only from its neighbors in a sequence [46] (see Figure 6-4a) or the communication topology may be more arbitrary [47] (see Figure 6-4b). Let us say that agents perform *partial public voting* when agents observe proper subsets of precedent local decisions.

Corollary 6.4. *Observing a subset of public signals does not affect optimal decision rules and performance of a team.*




Figure 6-4. Examples of partial public voting. (a) Each agent sequentially observes the decision made by the agent right before her. (b) Each agent observes the decisions made by her neighbors.

Proof. First, we will show that a team of agents observing only a subset of public signals (Team A) cannot outperform a team of the same size that consists of those who observe full public signals (Team B). The proof is by contradiction. Let us assume that Team A can outperform Team B using their optimal decision strategy. Then, since what each agent in Team A observes is also observed by the corresponding agents in Team B, Team B can mimic the optimal strategy of Team A. For Agent n in Team B, all she has to do is ignore the public signals that Agent n in Team A cannot observe. After mimicking the strategy of Team A, the performance of Team B becomes the same as that of Team A. This contradicts our assumption. Hence Team A cannot outperform Team B.

Next, let us consider Team A and a team of agents not observing any public signals (Team C). We can prove that Team C cannot outperform Team A through similar logic.

Corollary 6.3 implies that Team C in fact performs as well as Team B. Therefore Team A is also as good as Teams B and C with respect to their optimal performances.

□

The convenience of secret voting emerges especially when agents cannot observe all public signals. Even though partial public voting cannot outperform public voting, the former requires more computations for Bayesian social learning. When the agents observe a subset of public signals, they need to consider all possible realizations of the public signals that they cannot observe in order to perform Bayesian learning. For example, in Figure 6-4a, Carol observes Britta's decision but not Alexis's decision.

Her updated belief when $\widehat{H}_2 = 0$ will be computed as follows:

$$\begin{aligned}
q_3^{\text{A}0} &= \mathbb{P}\{H = 0 \mid \widehat{H}_2 = 0\} \\
&= \sum_{\widehat{h}_1} \mathbb{P}\{H = 0 \mid \widehat{H}_2 = 0, \widehat{H}_1 = \widehat{h}_1\} \mathbb{P}\{\widehat{H}_1 = \widehat{h}_1 \mid \widehat{H}_2 = 0\} \\
&= \sum_{\widehat{h}_1} \frac{\mathbb{P}\{H = 0 \mid \widehat{H}_2 = 0, \widehat{H}_1 = \widehat{h}_1\} \mathbb{P}\{\widehat{H}_1 = \widehat{h}_1, \widehat{H}_2 = 0\}}{\mathbb{P}\{\widehat{H}_1 = 0, \widehat{H}_2 = 0\} + \mathbb{P}\{\widehat{H}_1 = 1, \widehat{H}_2 = 0\}}. \tag{6.14}
\end{aligned}$$

This process is more complicated than belief update with the knowledge of both $\widehat{H}_1 = 0$ and $\widehat{H}_2 = 0$, which is just to compute $\mathbb{P}\{H = 0 \mid \widehat{H}_2 = 0, \widehat{H}_1 = 0\}$.

Instead of accepting this complexity, the agents should ignore the public signals. Since the optimal secret voting strategy performs equally to the optimal public voting strategy, it is economical for them to not share any public signals at all.

■ 6.3 Agents with Different Likelihoods

Agents may have private signals that relate differently to the hypothesis. For example, Agent n observes $Y_n = H + W_n$, where $W_n \sim \mathcal{N}(0, \sigma_n^2)$ is independent of but not identically distributed to the noise of Agent m , $W_m \sim \mathcal{N}(0, \sigma_m^2)$ for any $m \neq n$.² Signal-to-noise ratio (SNR) of the private signal Y_n is

$$(\text{SNR})_n = \frac{\mathbb{E}[H^2]}{\mathbb{E}[W_n^2]} = \frac{1 - p_0}{\sigma_n^2}.$$

Agent n observes a signal with relatively high signal-to-noise ratio (SNR) than Agent m does if $\sigma_n^2 < \sigma_m^2$. Like in Section 4.3, Agent n can be considered as an expert and Agent m as a novice. Their decisions are not equally informative, unlike in the identical-agent case of Section 6.1. It will be shown that the public signals are futile in cases where the fusion rule requires unanimity but useful in other cases.

Now that the agents' private signals do not have the same distributions, the order in which the agents act matters to their team performance. To distinguish the agents, we name them in descending order of SNRs: Amy has the highest SNR, Beth has the

²Even though additive Gaussian noises are considered in the discussion for convenience, the results hold for other kinds of private-signal model.

second-highest SNR, and so on. However, we keep the notation that the numbering of agents (and hence the subscript indices) indicate the order of decision making, i.e., Agent 1 acts first, but may or may not be Amy.

■ 6.3.1 AND and OR Fusion Rules

The AND (N -out-of- N) and OR (1-out-of- N) rules have a common feature. The team decision requires unanimity of the agents of one type or the other: for a team decision of 1 under the AND rule, the agents must unanimously decide 1; and for a team decision of 0 under the OR rule, they must unanimously decide 0. We thus call these *unanimity fusion rules*. This characteristic gives a special result for these fusion rules.

Let us consider a team of two agents with the OR rule: Amy with a higher SNR and Beth with a lower SNR. From the discussion of the two-agent case in Section 6.1, Amy and Beth each has one degree of freedom regardless of their order of decision making. Since the decision of the second-acting agent is irrelevant when the first-acting agent chooses 1, the second-acting agent only needs one decision threshold just in case when the first-acting agent chooses 0. Suppose Amy makes her decision first with decision threshold λ_A , and Beth then makes her decision with decision threshold λ_B^0 , regardless of Amy's decision. Their minimum Bayes risk is

$$\min_{\lambda_A, \lambda_B^0} c_{10} p_0 \left(P_{e,A}^I + (1 - P_{e,A}^I) P_{e,B}^{I_0} \right) + c_{01} (1 - p_0) P_{e,A}^{II} P_{e,B}^{II_0}. \quad (6.15)$$

Now suppose that they switch their positions: Beth first makes her decision with decision threshold ρ_B , and Amy then makes her decision with decision threshold ρ_A^0 , regardless of Beth's decision. Their minimum Bayes risk is now given by

$$\begin{aligned} & \min_{\rho_B, \rho_A^0} c_{10} p_0 \left(P_{e,B}^I + (1 - P_{e,B}^I) P_{e,A}^{I_0} \right) + c_{01} (1 - p_0) P_{e,B}^{II} P_{e,A}^{II_0} \\ & = \min_{\rho_B, \rho_A^0} c_{10} p_0 \left(P_{e,A}^{I_0} + (1 - P_{e,A}^{I_0}) P_{e,B}^I \right) + c_{01} (1 - p_0) P_{e,A}^{II_0} P_{e,B}^{II}, \end{aligned} \quad (6.16)$$

where the terms are rearranged to have the same form as (6.15). Note that $P_{e,A}^I$, a

function of λ_A , and $P_{e,A}^{I_0}$, a function of ρ_{A0} , are the same error functions. Likewise, each pair of $P_{e,A}^{II}$ and $P_{e,A}^{II_0}$, $P_{e,B}^{I_0}$ and $P_{e,B}^I$, and $P_{e,B}^{II_0}$ and $P_{e,B}^{II}$ is the same error function. Therefore, the optimal decision thresholds are also the same:

$$\lambda_A^* = \rho_A^{0*} \quad \text{and} \quad \lambda_B^{0*} = \rho_B^*.$$

In conclusion, not only is the public signal useless but also the agents' decision-making order does not affect the optimal team decision. Their optimal strategy is just to adopt the decision thresholds λ_1^* and λ_2^* that minimize the Bayes risk:

$$\begin{aligned} (\lambda_1^*, \lambda_2^*) = \arg \min_{(\lambda_1, \lambda_2)} \{ & c_{10} p_0 (1 - (1 - P_{e,1}^I)(1 - P_{e,2}^I)) \\ & + c_{01} (1 - p_0) P_{e,1}^{II} P_{e,2}^{II} \}. \end{aligned} \quad (6.17)$$

We can extend this result to N agents as long as the fusion is performed by the OR rule or the AND rule.

Theorem 6.5. *For a unanimity fusion rule, secret voting is the optimal strategy even when agents observe private signals with different SNRs. Specifically, public signals are useless and the ordering of agents does not affect their optimal decision rules nor the resulting performance.*

Proof. For the OR rule, where the Bayes risk is given by

$$c_{10} p_0 \left(1 - \prod_{n=1}^N (1 - P_{e,n}^I) \right) + c_{01} (1 - p_0) \prod_{n=1}^N P_{e,n}^{II}, \quad (6.18)$$

each agent has a meaningful decision threshold only if all previous agents declare 0. Otherwise, decisions of the remaining agents are irrelevant. Thus, without loss of optimality, we can constrain that the agents optimize their decision thresholds λ_n for the case when $\widehat{H}_1 = \widehat{H}_2 = \dots = \widehat{H}_{n-1} = 0$ and use the same decision threshold for any public signals. In fact, they need not know the public signals; they just need to perform decision making as if all public signals are 0. If this assumption is not true, i.e., any earlier-acting agent chooses 1, then their decisions will be irrelevant anyway.

Furthermore, the symmetry in (6.18) implies that indices of Agents m and n are interchangeable. Therefore, the ordering of agents does not affect the optimal values of decision thresholds and, consequently, the minimum Bayes risk.

For the AND rule, the Bayes risk is given by

$$c_{10}p_0 \prod_{n=1}^N P_{e,n}^I + c_{01}(1-p_0) \left(1 - \prod_{n=1}^N (1 - P_{e,n}^{II}) \right), \quad (6.19)$$

and we can prove the statement in a similar way. \square

In general, the solution to minimizing (6.18) or (6.19) does not satisfy $\lambda_1^* = \lambda_2^* = \dots = \lambda_N^*$, unlike in Section 6.1. Agents will use different optimal decision thresholds corresponding to their SNRs. As discussed in Section 4.3, they need to solve equations like (4.11) to get the optimal decision thresholds. However, they do not know other agents' SNRs, specifically, other agents' noise variances σ_n^2 . Therefore, their rational decision thresholds, which are obtained from equations like (4.12) or (4.13), are suboptimal. As in Section 4.3, a team of lucky imperfect agents can outperform the team of perfect agents in the case when the agents observe private signals that are not identically distributed.

■ 6.3.2 Other Fusion Rules

Optimal decision making is more complex for other fusion rules due to the increase of degrees of freedom. Even in the simplest case when three agents make a decision with the MAJORITY (2-out-of-3) rule, the last two agents have two meaningful degrees of freedom each. The second agent has different decision thresholds λ_2^0 and λ_2^1 for $\widehat{H}_1 = 0$ and for $\widehat{H}_1 = 1$, respectively, and the third agent has different ones λ_3^{01} and λ_3^{10} for $(\widehat{H}_1, \widehat{H}_2) = (0, 1)$ and for $(\widehat{H}_1, \widehat{H}_2) = (1, 0)$, respectively. The third agent is irrelevant for $(\widehat{H}_1, \widehat{H}_2) = (0, 0)$ or $(\widehat{H}_1, \widehat{H}_2) = (1, 1)$ because the team decision has been made without her decision. Learning from public signals can be helpful in decision making due to these extra degrees of freedom, unlike in Section 6.3.1

Our symmetric fusion rule always treats all local decisions with equal weights even though they are made by agents that experience different SNRs. Thus, the team




Figure 6-5. Lower bounds of operating regions for different orderings of three agents, Amy ($\sigma_A^2 = 0.25$), Beth ($\sigma_B^2 = 1$), and Cindy ($\sigma_C^2 = 2.25$), and the MAJORITY fusion rule.

decision can be improved by social learning, which inevitably unbalances weights of local decisions.

As evidence of helpful social learning, Figures 6-5 and 6-6 are presented. In the case of Figure 6-5, there are three agents using the MAJORITY fusion rule. Amy has the highest SNR ($\sigma_A^2 = 0.25$), Beth has the median SNR ($\sigma_B^2 = 1$), and Cindy has the lowest SNR ($\sigma_C^2 = 2.25$). Figure 6-5 depicts the optimal reversed ROC curves for all possible orderings of the actions of the three agents.³

Note that the three orderings presented are sufficient because the order of the last two agents does not matter. Once the first agent chooses 0, the fusion rule is changed to 2-out-of-2 (AND) rule for the other two agents. If the first agent chooses 1, the fusion rule is changed to 1-out-of-2 (OR) rule. As discussed in Section 6.3.1, the order of agents is irrelevant if the fusion rule is a unanimity rule.

There are two notable things in Figure 6-5. First, the reversed ROC curve of secret voting is above that of public voting for any ordering. This implies that public voting strictly outperforms secret voting, regardless of the order in which the agents make

³As decision threshold parameters are varied with the order of agent actions fixed, some set of (P_E^I, P_E^{II}) pairs is achievable. We call the lower boundary of this set the reversed ROC curve.




Figure 6-6. Lower bounds of operating regions for different orderings of four agents, Amy ($\sigma^2 = 0.25$), Beth ($\sigma^2 = 0.5$), Cindy ($\sigma^2 = 1$), and Daphne ($\sigma^2 = 2.25$), and the 2-out-of-4 fusion rule.

decisions. Second, among the public voting scenarios, the best is when Beth makes her decision first. Further numerical experiments for $N = 3$ with several different noise variances also show that the team performance is the best when the agent with the median SNR acts first.

Figure 6-6 shows the optimal reversed ROC curves for four agents and the 2-out-of-4 fusion rule. Again, there is no need to compare all 16 orderings to find the best-performing one. After the first agent makes a decision, the updated fusion rule will be the 1-out-of-3 rule if the decision is 1, or 2-out-of-3 rule if the decision is 0. Since the ordering of the next three agents is irrelevant under 1-out-of-3 (OR) rule, their optimal order under 2-out-of-3 rule should be considered. In Figure 6-6, considered are four orderings with different first agent and optimally arranged three other agents. It is shown that public voting always outperforms secret voting. In addition, the agent with the second-highest SNR should act first but the difference between the case when the agent with the third-highest SNR acts first is very small.

In this section, we have provided evidence that agents can exploit social learning to improve their team decision when the qualities of the private signals vary. The sequence of the agents also needs to be carefully chosen to achieve the best team

performance.

■ 6.4 Experimental Investigation of Human Social Learning in Team Decision Making

Corollary 6.3 is a notable result because individually making the best possible decisions with all available information may seem to be also the best policy for the team, but it is not actually fully rational. To test whether people react to public signals, we conducted experiments considering conditionally iid private signals and complete public signals as in Section 6.1.

The objective of these experiments is verification of the two following hypotheses:

Hypothesis 6.6. *Human decision makers are affected by public signals in team decision making.*

Hypothesis 6.7. *Human decision makers are less affected by public signals in team decision making than they are in individualistic decision making.*

Our first hypothesis argues that humans are not rational in their use of public signals; thus, to improve team performance, the public signals should be eliminated, i.e., votes should be kept private. Our second hypothesis argues that human behaviors are not completely irrational; having less dependence on public signals at least matches the trend of fully rational behavior.

■ 6.4.1 Experiment A: Setup

The experiment asks subjects to perform decision-making tasks in the scenario described below.⁴

1. Each subject is told they are one of seven contestants in a game show. The contestant is assigned a number from 1 to 7. This will be their decision-making order.

⁴Before starting the experiment, the author and the PI had passed a training course on human subjects research, which is required according to MIT regulation. They applied for exempt status for this experiment and got approved by the Committee on the Use of Humans as Experimental Subjects (COUHES): COUHES Protocol No. 1310005949.




Figure 6-7. The likelihoods of a private signal.

2. A basket is presented on the stage. The contestants are asked to guess whether it is Basket A, which has 70 red balls and 30 blue balls, or Basket B, which has 30 red balls and 70 blue balls.
3. Each contestant can draw 9 balls from the basket (as a private signal) and put them back so that other contestants cannot see what he/she draws. Of course, exchange of the information is not allowed.
4. From contestant 1 through 7, they speak out what they think the basket is. They can also hear other contestants' answers (as public signals).

This game corresponds to the following hypothesis testing problem. The hypothesis is $H = 1$ if the basket is Basket A and $H = 0$ if Basket B. The prior probability of each hypothesis is 0.5. The private signal Y_n is the number of red balls among the 9 balls. Its likelihood functions are given by

$$p_{Y_n|H}(y|0) = \frac{\binom{30}{y}\binom{70}{9-y}}{\binom{100}{9}} \quad \text{and} \quad p_{Y_n|H}(y|1) = \frac{\binom{70}{y}\binom{30}{9-y}}{\binom{100}{9}}, \quad (6.20)$$

which are depicted in Figure 6-7. Seven agents observe conditionally iid private signals and public signals. Without public signals, the rational behavior is to choose Basket A if the agent draws 5 or more red balls and to choose Basket B otherwise.

The experiment consists of two versions of questionnaires: one for a sequential detection task as in Chapter 5 and the other for a distributed detection task as in Chapter 6. In the former questionnaire, each contestant gets a reward if his/her guess is correct and gets nothing otherwise, which is an individualistic goal. In the latter questionnaire, all contestants get equal reward if the majority guesses the correct basket and get nothing otherwise, which is a team goal. Hypothesis 6.7 can be verified by comparing results of these two versions.

The subjects are 200 Amazon Mechanical Turk workers living in the United States. Each subject is randomly assigned one of the two versions so that each version is assigned 100 workers. The subjects are asked to answer their thresholds to choose Basket A for each of 29 combinations of public signals. They are not informed that there are two versions of questionnaires because awareness of the other version may force them to think that they should behave differently.

In addition, we offer each subject a bonus reward up to \$1.50 in order to encourage the subject to think deeply and answer sincerely, while the base reward for completing the experiment is \$0.50. After reviewing each subject's answers, we gave a bonus proportional to the expectation of the individual's (Version I) or the team's (Version II) correctness.

■ 6.4.2 Experiment A: Results

The subjects provide their answers with a slider bar with effective range of -0.5 to 9.5 and resolution 0.1. The answers are quantized with resolution 1.0 because the thresholds represent the minimum number of red balls to choose Basket A. For example, decision thresholds 4.1 and 4.9 yield the same decision rules because the private signals are integer-valued. Hence, for any integer $n \in \{0, 1, \dots, 10\}$, answers within $(n - 1, n]$ are quantized to $n - 0.5$. Then we take the average of the quantized answers.

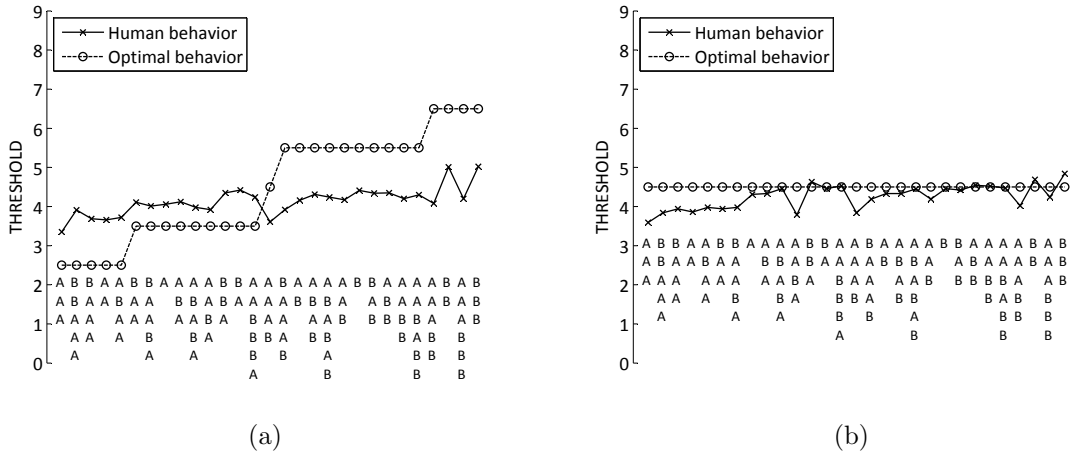


Figure 6-8. Optimal decision-making thresholds and empirical thresholds for various public signals. Each set of public signals shown to the subjects is written vertically, from top to bottom in the chronological order of the public signals. (a) Sequential detection—individualistic decision making (Version I). (b) Distributed detection—team decision making (Version II).

Figure 6-8a shows the results in the sequential detection task (Version I). This is individualistic decision making; contestants should do social learning to increase the probability of choosing the correct basket. The optimal threshold varies from 2.5 to 6.5 depending on public signals. The results reveal that human subjects do social learning but do not rely on public signal as much as they should do.

Figure 6-8b shows the results in the distributed detection task (Version II). This is team decision making so contestants should ignore public signals. Thus, the optimal decision threshold is 4.5 regardless of the public signals. However, Figure 6-8b clearly shows that humans are affected by public signals. Among 100 subjects who were assigned this version, only 3 subjects made their decisions independently of the public signals. We conclude that Hypothesis 6.6 is verified.

Figure 6-9 compares human behaviors in individualistic and in team decision-making task. The lengths and directions of arrows show how behavior is and should be changed from Version I to Version II.

Compared to the results of Version I, human decision makers do not seem to change their behaviors for individualistic or team goals. There are slight changes of the thresholds but they are too small for us to argue that the agents are behaving differently in the two cases. We ran a two-sample Kolmogorov-Smirnov test to compare




Figure 6-9. Comparison of decision thresholds in Version I and Version II. The directions of arrows indicate behavior changes from the individualistic case (Version I) to the team-objective case (Version II). Black arrows start from empirical thresholds for individualistic decision making to empirical thresholds for team decision making. For comparison, white arrows start from optimal thresholds for individualistic decision making to optimal thresholds for team decision making. The latter are all at 4.5 since the public signals are optimally ignored.

the distributions of empirical thresholds for each situation. In 25 out of 29 situations, they are the same distributions with probability higher than 50%. The probability is even higher than 90% in 14 situations. We conclude that Hypothesis 6.7 is not supported. More details are provided in Section 6.A.

Besides, an interesting behavior is observed when the public signals are [A A A B] in Figure 6-8a in Version I. Humans tend to use 3.5 as their thresholds even though the optimal threshold is 4.5. A rational interpretation of the public signals is that Agent 4 chooses Basket B because the agent has very strong signal for Basket B, i.e., the agent must have picked 7 or more blue balls. The likelihood of the event is negligibly small if the basket is Basket A according to Figure 6-7. Thus, subjects should take the decision of Agent 4 much more seriously than the decisions of the first three agents. That is why the optimal threshold is 4.5 even though 3 out of 4 public signals are Basket A.

However, 3.5 is the average threshold that human decision makers used. We interpret this behavior in two ways. First, humans may not believe that other people

are rational. They may think that Agent 4 is just wrong. Second, humans may process the public signals altogether regardless of the order of decisions made. Compared to other scenarios when they observe public signals such as [B A A A], [A B A A], [A A B A], their average decision thresholds are close to each other.

In conclusion, humans perform social learning even when they ought not to; they tend to use all available information. Furthermore, they do not seem to depend less on public signals for a team goal than they do for an individualistic goal.

■ 6.4.3 Experiment B: Setting and Results

The first experiment, Experiment A, considers the situation when each contestant draws nine balls from the basket. Letting them draw that many balls helps us survey subjects' decision-making thresholds at a fine scale. However, the downside is that the likelihood functions (6.20) are complicated. The subjects might use very similar thresholds in Versions I and II because the likelihood functions are too complicated for them to understand correctly.

We conducted another experiment⁵ with the same setting except simple private signals—each contestant draws only one ball from the basket. Contestants have only two cases: The ball they draw is red or blue. The likelihoods of their private signals Y_n are given by

$$\begin{cases} p_{Y_n|H}(\text{“red”} | 0) = 0.3, \\ p_{Y_n|H}(\text{“blue”} | 0) = 0.7, \end{cases} \quad \text{and} \quad \begin{cases} p_{Y_n|H}(\text{“red”} | 1) = 0.7, \\ p_{Y_n|H}(\text{“blue”} | 1) = 0.3, \end{cases}$$

where, again, $H = 1$ means that the basket is Basket A and $H = 0$ means Basket B.

In Experiment B, we discarded unreasonable combinations of public signals in the sense that herding occurs after one choice outnumbers the other by two votes in the public signals. For example, [A B A A B] would not occur because if the fifth contestant, who observe the public signals [A B A A], were rational, she would choose A regardless of her private signal as discussed in Section 2.3.2. A total of

⁵The second experiment has also been approved by the COUHES: COUHES Protocol No. 1403006285.

18 combinations of public signals were selected and asked. Each subject is asked to choose one basket when they draw a red ball and when they draw a blue ball. They are 100 Amazon Mechanical Turk workers who live in the United States and did not participate in Experiment A. Like in Experiment A, a half of the subjects was assigned to Version I and the other half to Version II.

Their decision making was categorized into four groups:

- *Not herding* (NH): to choose Basket A if the drawn ball is red and Basket B if it is blue.
- *Herding toward A* (HA): to choose Basket A regardless of the color of the drawn ball.
- *Herding toward B* (HB): to choose Basket B regardless of the color of the drawn ball.
- *Unreasonable choice* (UC): to choose Basket A if the drawn ball is blue and Basket B if it is red.

Since we considered the public signals that could only lead to NH or HA ideally, HB or UC should not appear if all subjects think rationally. However, a few subjects chose HB or UC as shown in Section 6.B. Especially, UC literally does not make sense because it is to choose the basket in which the drawn color is minority. We still respect such answers because the subjects were tested if they understood the given experiment setting and the rules of the game show; only those who answered correctly to all the questions were qualified to take the experiment.

Figure 6-10 shows decision making rules adopted by most subjects. Only two categories are displayed because most subjects adopted NH or HA. Figure 6-10a shows the results in the sequential detection task for individual correctness (Version I). It turns out that what most subjects did is optimal decision-making. Compared to Figure 6-8a, subjects were capable of thinking rationally with simple private signals like binary-valued ones.

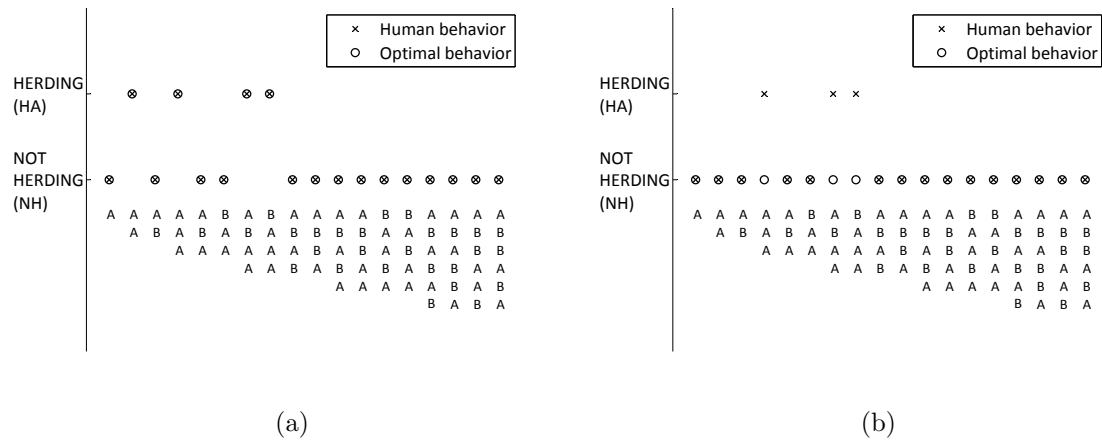


Figure 6-10. The behaviors shown by most subjects and their optimal behaviors. Each set of public signals shown to the subjects is written vertically, from top to bottom in the chronological order of the public signals. (a) Sequential detection—individualistic decision making (Version I). (b) Distributed detection—team decision making (Version II).

Figure 6-10b shows the results in the distributed detection task for team’s correctness (Version II). The subjects still did social learning. In fact, they made decisions which would have been rational if they had pursued individual correctness. The comparison of their decision rules in Versions I and II in Section 6.B supports the argument.

What can be learned from these two experiments is that many people have the ability to process public signal and perform social learning to make better individual decisions. Their cognitive power is limited to process complex private signals but still close to optimum when the private signals are simple. However, they lack understanding of how to make decisions to optimize the team’s correctness. They make decisions in a very similar way for themselves or for their teams irrespective of the simplicity of the private signals. Thus, if possible, it is better to not show irrelevant information that may confuse a person’s decision making.

■ 6.5 Conclusion

A combined model of a distributed detection and social learning has been discussed. Chapters 4 and 5 have separately discussed a distributed detection with symmetric

fusion and a sequential detection with social learning. Both techniques were used to decrease the Bayes risk. However, when they are combined, it is not a trivial question whether they are synergic.

The fundamental role of social learning is belief update, as discussed in Chapters 2 and 5. When Agent n chooses a hypothesis \hat{H}_n , the choice is not independent of H . When other agents watch the choice as a public signal, it is a rational and natural phenomenon that the public signal changes their beliefs about H . This is the only effect of social learning in a sequential detection scenario.

When social learning is performed within a decision-making team that aggregates opinions by voting, however, another effect of social learning arises. Because the public signals are the votes themselves cast by agents, the agents know the exact numbers of votes 0 or 1 required to reach the global decision 0 or 1 at every moment. They know the changing impact of a vote of 0 or 1 according to the evolution of the fusion rule. This fusion rule update is a less frequently explored topic.

When the agents observe conditionally iid private signals, the effects of the belief update and the fusion rule evolution cancel exactly. Consequently, the optimal performance with or without public signals is the same; internally flowing information does not improve the team performance.

When the agents observe signals that are conditionally independent but not identically distributed, social learning may improve team performance, and when it does, the degree of improvement depends on the order in which the agents act. Among the symmetric fusion rules, the 1-out-of- N and N -out-of- N fusion rules are peculiar in that they essentially require unanimity among agents. Social learning becomes meaningless because the unanimity rules prevent the agents from specializing their decision thresholds for various public signals.

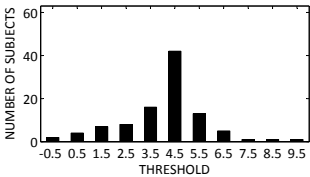
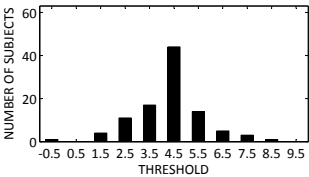
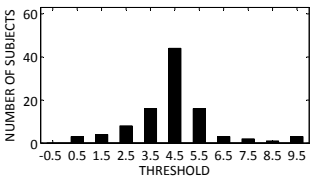
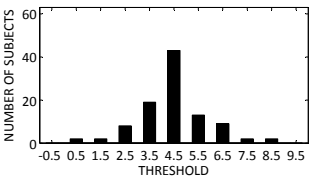
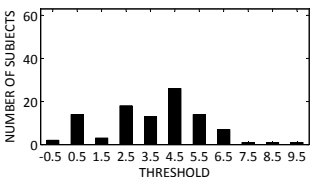
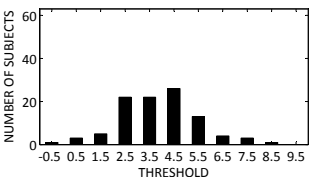
Social learning can play a role to improve team decisions as long as the fusion rule is not one of these unanimity rules. In examples with Gaussian likelihoods, we showed that team performance improves when agents with differing observation SNRs do social learning. With three agents making a team decision by the MAJORITY fusion rule, it is best for the agent with median SNR to act first. With four agents using

the 2-out-of-4 fusion rule, it is best for the agent with second-best SNR to act first.

These results do not argue that public signals make team performance worse. Even if the team of agents do social learning, they achieve the best team performance as long as they do it properly. Social learning causes a trouble in our experiments because humans overlook the evolution of fusion rule and act as if they are making decisions for individualistic goals. This emphasizes the importance of design of rules so that they can supplement humans' suboptimal behaviors.



■ 6.A Results of Experiment A

Table 6.1: Subjects' Decision Thresholds When They Draw Nine Balls

	Distribution of Threshold for the Individualistic Goal (Version I)	Distribution of Threshold for the Team Goal (Version II)	Two-Sample Kolmogorov- Smirnov Test ⁶
A			Accept H_0 with $p = 0.8938$
B			Accept H_0 with $p = 1.0000$
AA			Accept H_0 with $p = 0.4431$



⁶The null hypothesis H_0 is that the two set of samples are drawn from the same distribution.

AB





Accept H_0 with
 $p = 0.7942$

BA





Accept H_0 with
 $p = 1.0000$

BB





Accept H_0 with
 $p = 0.4431$

AAA





Accept H_0 with
 $p = 0.7942$

AAB





Accept H_0 with
 $p = 0.9995$

ABA





Accept H_0 with
 $p = 0.1930$

ABB





Accept H_0 with
 $p = 0.9995$

BAA





Accept H_0 with
 $p = 0.9921$

BAB





Accept H_0 with
 $p = 1.0000$

BBA





Accept H_0 with
 $p = 0.8938$

BBB





Accept H_0 with
 $p = 0.6766$

AAAB





Accept H_0 with
 $p = 0.9610$

AABA





Accept H_0 with
 $p = 0.9921$

ABAA





Accept H_0 with
 $p = 0.8938$

BAAA





Accept H_0 with
 $p = 0.8938$

AABB





Accept H_0 with
 $p = 0.6766$

ABAB





Accept H_0 with
 $p = 1.0000$

AAABB





Accept H_0 with
 $p = 0.9995$

ABABA





Reject H_0 with
 $p = 0.0314$

BBAAA

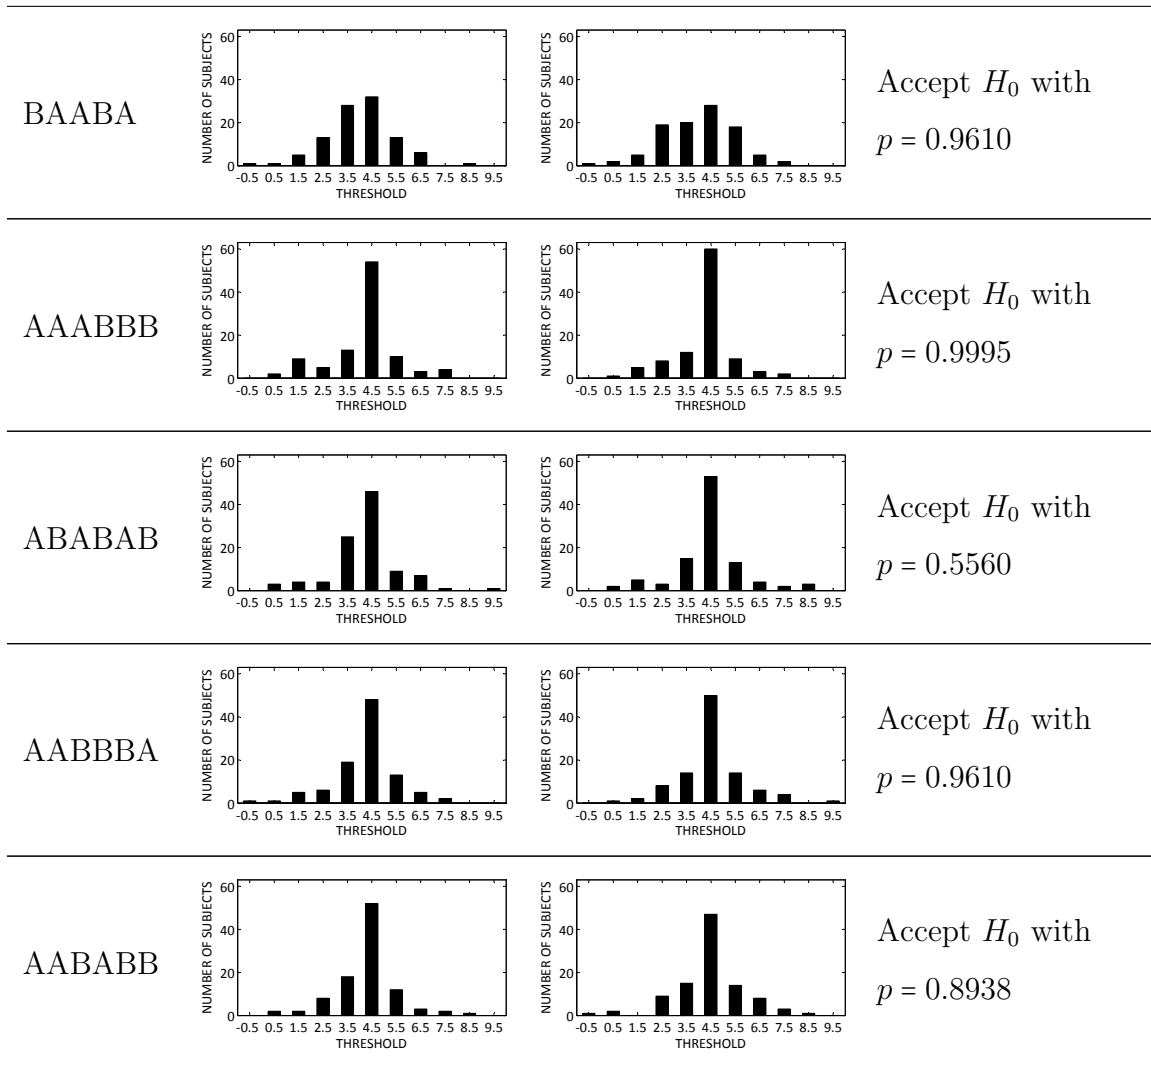


Accept H_0 with
 $p = 1.0000$

BAAAB



Accept H_0 with
 $p = 0.0994$



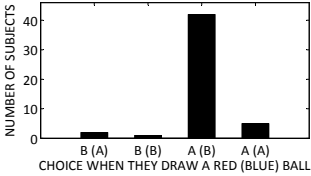
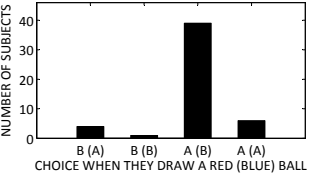
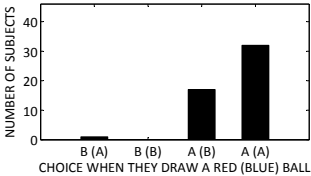
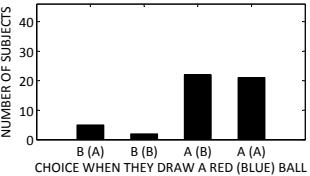
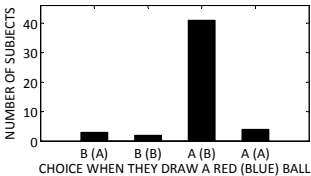
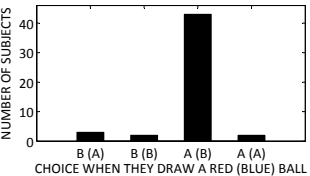
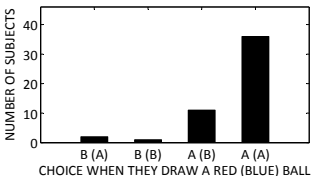
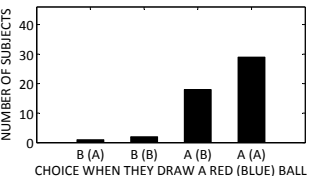
■ 6.B Results of experiment B

Table 6.2 shows how subjects made decisions when they drew a red or a blue ball. There are four possible choices. Among them, two kinds of herding occur: Always choosing Basket A or always choosing Basket B. Even though this experiment does not consider any public signal that optimally leads to herding to Basket B, some subjects chose Basket B regardless of their drawing. In addition, subjects who respect their private signals choose Basket A when they draw a red ball and Basket B otherwise. The opposite behavior—choosing Basket B when they draw a red ball and Basket



A otherwise—was also observed. This behavior does not make sense but we do not exclude such answers because they are not significant and we already have a qualification step to check whether the subjects fully understood the given experiment setting.

The two-sample Kolmogorov-Smirnov test accepts the null hypothesis for all public signal cases. It implies that subjects could not determine the optimal strategy for team decision-making goals even when their private signals are so simple that they only need to consider two cases—what to choose when they draw a red ball and what to choose when they draw a blue ball.

Table 6.2: Subjects' Decision Thresholds When They Draw One Ball



	Distribution of Threshold for the Individualistic Goal (Version I)	Distribution of Threshold for the Team Goal (Version II)	Two-Sample Kolmogorov- Smirnov Test
A			Accept H_0 with $p = 1.0000$
AA			Accept H_0 with $p = 0.1546$
AB			Accept H_0 with $p = 1.0000$
AAA			Accept H_0 with $p = 0.6779$

ABA





Accept H_0 with
 $p = 1.0000$

BAA





Accept H_0 with
 $p = 0.8409$

ABAA





Accept H_0 with
 $p = 1.0000$

BAAA





Accept H_0 with
 $p = 0.9541$

ABAB





Accept H_0 with
 $p = 1.0000$

ABBA





Accept H_0 with
 $p = 0.9541$

ABABA





Accept H_0 with
 $p = 1.0000$

ABBAA





Accept H_0 with
 $p = 0.8409$

BAABA





Accept H_0 with
 $p = 0.9958$

BABAA





Accept H_0 with
 $p = 1.0000$

ABABAB





Accept H_0 with
 $p = 1.0000$

ABABBA





Accept H_0 with
 $p = 1.0000$

ABBAAB



Accept H_0 with
 $p = 1.0000$

ABBABA



Accept H_0 with
 $p = 1.0000$

Conclusion

Perfect Bayesian decision making is computationally challenging. It requires complete understanding of likelihoods of hypotheses and observations and processing of them to get a posteriori probabilities. Some systems like humans have limited resources to perform a decision-making task. Among many difficulties of human decision makers in practical decision-making situations, the thesis focuses on their cognitive limitations toward the prior probability.

This thesis introduces an imperfect decision-making agent, which models human cognitive processes. Psychological experiments have revealed that human decision making relies on the appearance of the subjects, such as criminal defendants, students, and job applicants. Humans not only are more generous to attractive people but also use the attractiveness as a clue to perceive the prior probability. Inspired by the latter, the imperfect agent is assumed to estimate the prior probability upon observing an extra signal correlated to the prior probability. The Bayes-optimal estimation is the expectation of the prior probability conditioned on the signal, but this is relaxed so that the agent estimates the prior probability in an arbitrary way to an arbitrary value. In the discussions of Chapters 4, 5, and 6, imperfect agents who have arbitrary beliefs and perfect agents who have the correct prior probability are considered.

The main objective of this thesis is to investigate the performance of imperfect agents in terms of Bayes risk from the perspective of costly rationality—the agents optimize their behaviors with their limited abilities and resources. The relation of incorrect belief to the Bayes risk in a single-agent case is trivial. The agent will perform a suboptimal likelihood ratio test that is designed based on her incorrect

belief so she will have a higher Bayes risk. The relation is nontrivial in multiple-agent cases because of interactions between imperfect agents who have different beliefs. The thesis contributes to understanding fundamental principles of group decision making by looking at simple decision-making models containing multiple agents.

Incorrect beliefs are detrimental in distributed detection with voting and symmetric fusion when all agents observe conditionally iid private signals. On the other hand, they can be beneficial when the agents observe conditionally independent but not identically distributed signals. Perfect agents would use suboptimal decision rules in that case but imperfect agents have a chance to have better decision rules according to their beliefs.

Notable results discussed in the thesis include that some incorrect beliefs are beneficial to boost the efficiency of social learning in a sequential detection model. In addition, social learning can be futile in the distributed detection scenario if all agents observe conditionally iid private signals or the fusion rule is unanimous.

The main results of the thesis are further detailed in Section 7.1. Some possible future works inspired by the thesis are mentioned in Section 7.2.

■ 7.1 Recapitulation

The results can be classified with respect to the decision-making goal: a team goal as in the distributed detection model and an individual goal as in the sequential detection model. For these two different goals, the effects of incorrect beliefs and social learning are separately summarized below.

Having Incorrect Beliefs in Distributed Detection and in Sequential Detection

When agents are in a team that shares the cost of decision making, incorrect beliefs increase the team Bayes risk if the agents observe conditionally iid private signals. They can decrease the Bayes risk by gathering more agents to the team but cannot outperform a perfect team with the same number of agents whether the imperfect team consists of identical agents or diverse agents. Forming a diverse team has only an advantage with respect to stability.

The imperfect team can outperform the team of perfect agents only when the agents observe conditionally independent but not identically distributed private signals. The perfect agents cannot have optimal decision thresholds because they do not know the likelihood functions of other agents' private signals. Compared to them, the imperfect agents can have either better or worse decision rules according to their beliefs.

On the other hand, when the agents make decisions sequentially and individually, having correct beliefs causes suboptimal decision making even if the agents observe conditionally iid private signals. Correct belief helps an agent make a right decision but the decision is not as informative as a decision based on a slightly open-minded belief. Therefore, from the perspective of the last agent, who treats all earlier-acting agents as her advisors, it is better for the earlier-acting agents to be open-minded.

Social learning in Distributed Detection and in Sequential Detection

When agents make decisions individually, social learning plays a role of belief update. An agent who observes other agents' decisions as public signals can update her belief to make a better choice. In addition, it can lead to a better result if the public signals are generated by open-minded agents.

Social learning also plays the same role of belief update when agents make a decision as a team. However, it has another role in the case when the team decision is aggregated by voting with symmetric fusion: Social learning enables the agents to catch up with the evolution of the fusion rule. The fusion rule evolution countervails the belief update; they exactly cancel when the agents observe conditionally iid private signal.

Even though social learning does not increase Bayes risk, it may distract agents from optimal decision making. The experiment on human decision making showed that people do update their beliefs but do not properly incorporate the change of fusion rule. Therefore, it is economical and more efficient not to show them any public signals when they are useless.

Finally, it is shown that social learning can improve team decisions when the agents

observe conditionally independent but not identically distributed private signals. The degree of improvement depends on the order in which the agents make their decisions. For example, with three agents making a team decision by the MAJORITY fusion rule, it is best for the agent with median SNR to act first. With four agents using the 2-out-of-4 fusion rule, it is best for the agent with second-best SNR to act first.

■ 7.2 Future Directions

Theoretical Research on Decision Making

One of the most interesting observations is that the order in which agents make their decisions matters when the agents perform distributed detection with social learning and observe not identically distributed private signals. The order is not trivial and it has not been revealed yet why it is the best for the agent with median SNR to act first. Study of the optimal ordering will help us understand dynamics of agents in the team and especially the roles of the first agent beyond that being open-minded is related to being informative.

Practical Research on Decision Making

The experiment in Section 6.4 revealed an interesting irrational human behavior when humans absorb a public signal overturning overwhelming decisions. Social learning from the public signals [A A A B] yielded almost the same results as social learning from [B A A A] even though their implications are completely different. The former means that Agent 4 has a very strong private signal supporting B while the latter can occur when the private signal of Agent 1 only weakly supports B. Therefore, a rational decision maker should take the public signal B much more seriously in the former case. We do not know yet exactly why people do not take the signal rationally. It may be because they do not trust other people or because they do not consider the chronological order of public signals.

More generally, it needs to be studied how people process public signals and update their beliefs. Combined with the theoretical research, such research could be

influential widely on practical applications of social learning like political campaigns, marketing, and the advertisement business.

Bibliography

- [1] J. A. Swets, W. P. Tanner, Jr., and T. G. Birdsall, "Decision processes in perception," *Psychological Review*, vol. 68, no. 5, pp. 301–340, Sep. 1961. 16, 42
- [2] W. K. Viscusi, "Are individuals Bayesian decision makers?" *American Economics Review*, vol. 75, no. 2, pp. 381–385, May 1985. 16, 42
- [3] G. L. Brase, L. Cosmides, and J. Tooby, "Individuation, counting, and statistical inference: The role of frequency and whole-object representations in judgment under uncertainty," *Journal of Experimental Psychology: General*, vol. 127, no. 1, pp. 3–21, Mar. 1998. 16, 42
- [4] M. Glanzer, A. Hilford, and L. T. Maloney, "Likelihood ratio decisions in memory: Three implied regularities," *Psychonomic Bulletin & Review*, vol. 16, no. 3, pp. 431–455, Jun. 2009. 16, 42
- [5] B. W. Darby and D. Jeffers, "The effects of defendant and juror attractiveness on simulated courtroom trial decisions," *Social Behavior and Personality*, vol. 16, no. 1, pp. 39–50, Jun. 1988. 16, 44
- [6] M. J. Brown, E. Henriquez, and J. Groscup, "The effects of eyeglasses and race on juror decisions involving a violent crime," *American Journal of Forensic Psychology*, vol. 26, no. 2, pp. 25–43, Sep. 2008. 16, 45, 48
- [7] R. Radner, "Costly and bounded rationality in individual and team decision-making," *Industrial and Corporate Change*, vol. 9, no. 4, pp. 623–658, Dec. 2000. 17, 42

- [8] K. R. Varshney and L. R. Varshney, “Quantization of prior probabilities for hypothesis testing,” *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 4553–4562, Oct. 2008. 17, 18
- [9] C. N. Macrae and G. V. Bodenhausen, “Social cognition: Thinking categorically about others,” *Annual Review of Psychology*, vol. 51, pp. 93–120, Feb. 2000. 18
- [10] J. B. Rhim, L. R. Varshney, and V. K. Goyal, “Collaboration in distributed hypothesis testing with quantized prior probabilities,” in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 2011, pp. 303–312. 18
- [11] —, “Benefits of collaboration and diversity in teams of categorically-thinking decision makers,” in *Proc. IEEE Sensor Array and Multichannel Signal Processing Workshop*, Hoboken, NJ, Jun. 2012. 18
- [12] —, “Quantization of prior probabilities for collaborative distributed hypothesis testing,” *IEEE Transactions on Signal Processing*, vol. 60, no. 9, pp. 4537–4550, Sep. 2012. 18, 49
- [13] —, “Distributed decision making by categorically-thinking agents,” in *Decision Making with Imperfect Decision Makers*, ser. Intelligent Systems Reference Library, M. Kárný, T. V. Guy, D. Wolpert, and A. E. P. Villa, Eds. Springer, 2013. 18
- [14] Z. Chair and P. K. Varshney, “Optimal data fusion in multiple sensor detection systems,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-22, no. 1, pp. 98–101, Jan. 1986. 29
- [15] J. N. Tsitsiklis, “Decentralized detection,” in *Advances in Statistical Signal Processing*, vol. 2. JAI Press, 1993, pp. 297–344. 29, 32
- [16] —, “Decentralized detection by a large number of sensors,” *Mathematics of Control, Signals, and Systems*, vol. 1, no. 2, pp. 167–182, 1988. 32, 52

- [17] A. V. Banerjee, “A simple model of herd behavior,” *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, Aug. 1992. 35, 79
- [18] S. Bikhchandani, D. Hirshleifer, and I. Welch, “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, Oct. 1992. 35, 79
- [19] L. Smith and P. Sørensen, “Pathological outcomes of observational learning,” *Econometrica*, vol. 68, no. 2, pp. 371–398, Mar. 2000. 35, 39
- [20] M. Abdellaoui and E. Kemel, “Eliciting prospect theory when consequences are measured in time units: “Time is not money”,” *Management Science*, pp. 1–16, Dec. 2013, article in advance. 41
- [21] D. E. Bell, “Regret in decision making under uncertainty,” *Operations Research*, vol. 30, no. 5, pp. 961–981, Sep. 1982. 41
- [22] G. Loomes and R. Sugden, “Regret theory: An alternative theory of rational choice under uncertainty,” *The Economic Journal*, vol. 92, no. 368, pp. 805–824, Dec. 1982. 41
- [23] D. E. Bell, “Disappointment in decision making under uncertainty,” *Operations Research*, vol. 33, no. 1, pp. 1–27, Jan. 1985. 41
- [24] C. Camerer and M. Weber, “Recent developments in modeling preferences: Uncertainty and ambiguity,” *Journal of Risk and Uncertainty*, vol. 5, no. 4, pp. 325–370, Oct. 1992. 41
- [25] M. Rabin, “Incorporating fairness into game theory and economics,” *The American Economic Review*, vol. 83, no. 5, pp. 1281–1302, Dec. 1993. 41, 84
- [26] H. A. Simon, “A behavioral model of rational choice,” *The Quarterly Journal of Economics*, vol. 69, no. 1, pp. 99–118, Feb. 1955. 41
- [27] —, “Rational choice and the structure of the environment,” *Psychological Review*, vol. 63, no. 2, pp. 129–138, Mar. 1956. 41

- [28] J. Marschak and R. Radner, *Economic Theory of Teams*. New Haven, CT: Yale University Press, 1972. 42
- [29] J. E. Stewart, "Defendant's attractiveness as a factor in the outcome of criminal trials: An observational study," *Journal of Applied Social Psychology*, vol. 10, no. 4, pp. 348–361, Aug. 1980. 42
- [30] M. G. Efran, "The effect of physical appearance on the judgment of guilt, interpersonal attraction, and severity of recommended punishment in a simulated jury task," *Journal of Research in Personality*, vol. 8, no. 1, pp. 45–54, Jun. 1974. 43
- [31] G. Leventhal and R. Krate, "Physical attractiveness and severity of sentencing," *Psychological Reports*, vol. 40, no. 1, pp. 315–318, Feb. 1977. 43
- [32] C. Stephan and J. C. Tully, "The influence of physical attractiveness of a plaintiff on the decisions of simulated jurors," *Journal of Social Psychology*, vol. 101, no. 1, pp. 149–150, Feb. 1977. 43
- [33] K. Dion, E. Berscheid, and E. Walster, "What is beautiful is good," *Journal of Personality and Social Psychology*, vol. 24, no. 3, pp. 285–290, Dec. 1972. 43
- [34] M. M. Clifford and E. Walster, "The effect of physical attractiveness on teacher expectations," *Sociology of Education*, vol. 45, no. 2, pp. 248–258, 1973. 44
- [35] W. Goldman and P. Lewis, "Beautiful is good: Evidence that the physically attractive are more socially skillful," *Journal of Experimental Social Psychology*, vol. 13, no. 2, pp. 125–130, Mar. 1977. 44
- [36] C. M. Marlowe, S. L. Schneider, and C. E. Nelson, "Gender and attractiveness biases in hiring decisions: Are more experienced managers less biased?" *Journal of Applied Psychology*, vol. 81, no. 1, pp. 11–21, Feb. 1996. 44
- [37] H. Sigall and N. Ostrove, "Beautiful but dangerous: Effects of offender attractiveness and nature of the crime on juridic judgment," *Journal of Personality and Social Psychology*, vol. 31, no. 3, pp. 410–414, Mar. 1975. 44

- [38] E. D. Smith and A. Hed, “Effects of offenders’ age and attractiveness on sentencing by mock juries,” *Psychological Reports*, vol. 44, no. 3, pp. 691–694, Jun. 1979. 44
- [39] H. A. David and H. N. Nagaraja, *Order Statistics*, 3rd ed. Hoboken, NJ: John Wiley & Sons, 2003. 49
- [40] D. Teichroew, “Tables of expected values of order statistics and products of order statistics for samples of size twenty and less from the normal distribution,” *Annals of Mathematical Statistics*, vol. 27, no. 2, pp. 410–426, Jun. 1956. 49
- [41] E. Fehr and K. M. Schmidt, “A theory of fairness, competition, and cooperation,” *The Quarterly Journal of Economics*, vol. 114, no. 3, pp. 817–868, Aug. 1999. 85
- [42] G. Charness and M. Robin, “Understanding social preferences with simple test,” *The Quarterly Journal of Economics*, vol. 117, no. 3, pp. 817–869, Aug. 2002. 86
- [43] M. Ottaviani and P. Sørensen, “Information aggregation in debate: Who should speak first?” *Journal of Public Economics*, vol. 81, no. 3, pp. 393–421, Sep. 2001. 95
- [44] V. Krishnamurthy and H. V. Poor, “Social learning and Bayesian games in multi-agent signal processing: How do local and global decision makers interact?” *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 43–57, May 2013. 95
- [45] R. Radner, “Team decision problems,” *Annals of Mathematical Statistics*, vol. 33, no. 3, pp. 857–881, Sep. 1962. 97
- [46] P. K. Varshney, *Distributed Detection and Data Fusion*. Springer-Verlag, 1997. 105
- [47] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, “Bayesian learning in social networks,” *Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, Mar. 2011. 105