Quality of Information, Survival, and Incentives

by

Kookyoung Han

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2018

Doctoral Committee:

Professor Tilman Börgers, Chair
Assistant Professor Heng Liu
Associate Professor David Miller
Professor Uday Rajan

Kookyoung Han

kyhan@umich.edu

ORCID iD: 0000-0001-6601-6548

Acknowledgments

First of all, I would like to appreciate my committee chair, Tilman Börgers for his expertise, advice, encouragement and patience throughout my Ph.D. program. Since he has been constantly supporting me, I can finish my dissertation. I would like to thank my committee members, Heng Liu, David Miller, and Uday Rajan for their comments and discussions on my dissertation.

I would like to give special thanks to Linda Tesar for her help. Also, I would like to give special thanks to Stephan Lauermann at Universität Bonn. I am grateful to the Department of Economics.

Finally, I thank my family, my friends, and all seminar participants.

Contents

Acknowledgments **ii**

List of Figures **v**

List of Appendices **vi**

Abstract **vii**

Chapter

1 Introduction **1**

2 Model Uncertainty, Self-Enforcement, and Long Run Survival **4**

2.1 Introduction 4

2.1.1 Related Literature 9

2.2 Example 14

2.3 Model: A Risk Sharing Game 16

2.4 Bayesian Players and External Enforcement 20

2.5 Bayesian Players and Self Enforcement 30

2.6 Further Comments on Bayesian Players and Self Enforcement 37

2.7 Ambiguity Averse Players and External Enforcement 38

2.8 Ambiguity Averse Players and Self Enforcement 43

2.9 Conclusion 47

3 Screening Agents with Different Qualities of Information 51

3.1 Introduction 51

3.1.1 Related Literature 57

3.2 Model 59

3.3 Proper Scoring Rules 67

3.4 Relationship Between Proper Scoring Rules and Incentive Compatible Contracts 70

3.5 Further Properties of Incentive Compatible Contracts 74

3.6 Optimal Contracts 75

3.7 Linear Contracts 84

3.8 Example 89

3.9 Discussion of Assumptions 1 and 2 94

3.10 Conclusion 96

4 On The Optimal Order of Experiments 99

4.1 Introduction 99

4.1.1 Related Literature 104

4.2 Model 105

4.3 Normalization of Payoffs and the Cost of Experimentation . . 108

4.4 Fixed Cost of Experimentation: First Class of Examples . . . 113

4.5 Fixed Cost of Experiments: Comparison of FG and GF . . . 118

4.6 Fixed Cost of Experimentation: Second Class of Examples . . 123

4.7 Delay by Experimentation: First Class of Examples 126

4.8 Delay by Experimentation: Comparison of FG and GF 130

4.9 Delay by Experimentation: Second Class of Examples 130

4.10 Conclusion 132

Appendices 135

Bibliography 174

List of Figures

3.1 Geometric presentation of ξ and $\hat{\lambda}$ 64

3.2 Graphical illustration of Lemma 3.5.2 76

3.3 Affine transformation that makes the participation constraint of the uninformed type binding. 79

3.4 A small affine transformation while the uninformed type remains indifferent between accept and rejecting 80

3.5 Affine transformation that does not alter the lower threshold . 83

3.6 Optimal linear contract 85

3.7 Comparison of the principal's expected utility, the expected payment to the agent and the agent's outside option: $\tilde{g}$ is the best affine transformation of g. 92

3.8 Expected net benefits to the principal under $\tilde{g}$ 93

3.9 Proper scoring rule that corresponds to g 93

3.10 Optimal linear contract under different assumptions 97

List of Appendices

Appendix A. Omitted Proofs in Chapter 2 136
Appendix B. Omitted Proofs in Chapter 3 145
Appendix C. Omitted Proofs in Chapter 4 155

Abstract

This dissertation is devoted to questions on long run survival, the optimal elicitation of private information, and the optimal order of gathering information.

In Chapter 2, I consider an infinite horizon risk sharing game in which players have heterogeneous priors about future endowments, and analyze asymptotic behavior of efficient allocation depending on whether the players have commitment power and whether the players are Bayesian or ambiguity averse (Gilboa and Schmeidler (1989)). As in Blume and Easley (2006), I show that if the players are expected utility maximizing Bayesian learners and have commitment power, only survivors are those with the least incorrect beliefs. All other players starve in the long run. In other cases, no player vanishes. When the players are Bayesian and have no commitment power, no player starves in a Pareto efficient subgame perfect equilibrium. When the players are ambiguity averse and have commitment power, they can agree on a stationary allocation, which means that no player vanishes. When the players are ambiguity averse and have no commitment power, for sufficiently large discount factors, a stationary Pareto efficient allocation with commitment is a subgame perfect equilibrium.

In Chapter 3, I consider a principal-agent problem in which a principal elicits an agent's information when the quality of information provided by the agent depends on the agent's type. We investigate the impact of the agent's type dependent outside option on the optimal contract. Under restrictive assumptions on the type dependent outside option and the agent's

information structure, I show that the principal admits bad types and good types, but reject intermediate types. By further restricting our attention to a smaller class of decision problems, I show the existence of an optimal contract and construct how to design an optimal contract. Finally, I provide an example in which the principal optimally hires bad types to reduce the expected payment to good types. In the example, the principal actually loses if the agent draws a bad type.

In Chapter 4, co-authored with Professor Tilman Börgers, we study the optimal order of experimentation, considering a class of dynamic decision problems in which two experiments are available and a decision maker incurs costs of experimentation. Given the class of two binary experiments, there is no non-trivial comparison of sequential experiments. The reason why the decision maker runs a less informative experiment first in some circumstances is because the less informative experiment triggers the second experiment less frequently than the more informative experiment does. This idea allows us to come up with another class of two experiments, for which there exists non-trivial comparison of experiments. Given the second class of experiment, informativeness of static decision problems implies informativeness of dynamic decision problems. That is, it is optimal for the decision maker to run a more informative experiment first in every decision problem under study.

Chapter 1

Introduction

In this dissertation, I address questions on long run survival, the optimal elicitation of private information, and the optimal order of gathering information. Regarding long run survival, the existing literature supports the idea that when agents have heterogeneous priors in complete markets, the only survivors are those who know the truth. This result is partially because every agent in a complete market believes for sure that she eventually wins and locks herself into a long term contract. However, it would be difficult to find long run commitment. Agents may not be Bayesian, and may want to insure themselves against model uncertainty. I investigate the impact of a lack of commitment and the impact of ambiguity aversion on asymptotic behaviors of Pareto efficient allocations.

For the optimal elicitation of private information, an agent's outside option is crucial in designing the optimal contract. The existing literature has been focusing on how to induce the agent to truthfully report his private information, assuming that the agent always participates. However, the agent may not participate because of his large outside option. I study the impact of the agent's type dependent outside option on the optimal contract.

The optimal order of gathering information has been an interesting topic in economics. When it comes to dynamic decision problems, a decision maker is interested in not only the optimal action but also the optimal strategy of collecting information. When the decision maker collects information today, she should consider not only the cost of collecting information, but also impacts of today information on her optimal action and future information. However, it would be difficult to characterize the optimal order in which the decision maker collects information for general dynamic decision problems. Professor Tilman Börgers and I consider a class of decision problems that have been widely studied in economics, and look for the possibility that the optimal order in which to collect information is independent of decision problems under study.

In Chapter 2, I consider an infinite horizon risk sharing game in which players have heterogeneous priors about future endowments, and analyze the asymptotic behavior of efficient allocation depending on whether the players have commitment power and whether the players are Bayesian or ambiguity averse (Gilboa and Schmeidler (1989)). As in Blume and Easley (2006), I show that if the players are expected utility maximizing Bayesian learners and have commitment power, only survivors are those with the least incorrect beliefs. All other players starve in the long run. In other cases, no player vanishes. When the players are Bayesian and have no commitment power, no player starves in a Pareto efficient subgame perfect equilibrium. When the players are ambiguity averse and have commitment power, they can agree on a stationary allocation, which means that no player vanishes. When the players are ambiguity averse and have no commitment power, for sufficiently large discount factors, a stationary Pareto efficient allocation with commitment is a subgame perfect equilibrium.

In Chapter 3, I consider a principal-agent problem in which a principal elicits an agent's information when the quality of information provided by the agent depends on the agent's type. I investigate the impact of the agent's type dependent outside option on the optimal contract. Under restrictive assumptions on the type dependent outside option and the agent's information structure, I show that the principal admits bad types and good types, but reject intermediate types. By further restricting our attention to a smaller class of decision problems, I show the existence of an optimal contract and construct how to design an optimal contract. Finally, I provide an example in which the principal optimally hires bad types to reduce the expected payment to good types. In the example, the principal actually loses if the agent draws a bad type.

In Chapter 4, co-authored with Professor Tilman Börgers, we study the optimal order of experimentation, considering a class of dynamic decision problems in which two experiments are available and a decision maker incurs the costs of experimentation. Given the class of two binary experiments, there is no non-trivial comparison of sequential experiments. The reason why the decision maker runs a less informative experiment first, in some circumstances, is because the less informative experiment triggers the second experiment less frequently than the more informative experiment does. This idea allows us to come up with another class of two experiments, for which there exists a non-trivial comparison of experiments. Given the second class of experiment, informativeness of static decision problems implies informativeness of dynamic decision problems. That is, it is optimal for the decision maker to run a more informative experiment first in every decision problem under study.

Chapter 2

Model Uncertainty, Self-Enforcement, and Long Run Survival

2.1 Introduction

If we take a look around, we can see that many people have different beliefs about the same subject. For instance, we can think of sports betting. Theories of standard Bayesian expected utility imply that people can expect some gain from betting when they have different priors. Here, we notice that commitment and a person's confidence in the assessment of her prior are necessary for betting to take place. Requirement of commitment is clear because otherwise an agent can default and not pay her bet. An agent's confidence about her prior affects her willingness to bet. We can imagine a person who hesitates to bet because she has not collected enough information to form a sharp prior.

Blume and Easley (2006) is a seminal paper regarding the long run behavior of efficient allocations in the presence of heterogeneous beliefs. It shows that in a stationary exchange economy with the identical discount factors, when markets are complete and the allocation is Pareto efficient, only agents with the least incorrect beliefs survive in the long run. All other agents starve in the long run.[1] However, this result seems less intuitive in the sense that it

[1] The main theorems in Blume and Easley (2006) deal with general stochastic processes.

would be difficult to find the long run commitment. It is necessary for complete markets because they require external enforcement of infinite horizon contracts. Put differently, it would be difficult to imagine that an agent can commit to transferring most of her future endowment streams to someone. Focusing on an independently and identically distributed economy with the identical discount factors, we raise the following question: is it true that the only survivors are those with the least incorrect beliefs in the absence of commitment power?

Ambiguity aversion would be another factor that prevents an agent from vanishing. Billot *et al.* (2000), Rigotti, Shannon and Strzalecki (2008) and Ghirardato and Siniscalchi (2016) study efficient allocations when agents behave according to the maxmin criterion proposed in Gilboa and Schmeidler (1989). In those papers, it is shown that in complete markets, if aggregate endowments are constant and all ambiguity averse agents share a common prior, Pareto efficiency implies that every agent's consumption is constant across states. That is, no agent starves. Condie (2008), Da Silva (2011) and Guerdjikova and Sciubba (2015) study conditions under which an ambiguity averse player survives in the presence of an expected utility maximizing Bayesian learner with the correct belief. To depart from Billot *el at.* (2000) and related papers, we allow for aggregate uncertainty. In our model, an ambiguity averse player does not compete against a Bayesian expected utility maximizer who knows the truth. Then, we study the long run behavior of Pareto efficient allocation with and without commitment when all players are ambiguity averse.

In this paper, we study the impact of self enforcement and the impact of ambiguity aversion on the asymptotic behavior of efficient outcomes in an infinite horizon risk sharing game with finite players. In the risk sharing game, a player's endowment in each period is determined by a state, which is independently and identically distributed across periods. To study the impact of self enforcement on efficient allocations, we compare Pareto efficient allocations under external enforcement and Pareto efficient subgame perfect equilibriums of the risk sharing game when the players are expected utility maximizing Bayesian learners. To study the impact of ambiguity aversion, we analyze the efficient outcomes of complete markets when all players are Gilboa-Schmeidler (1989) type and aggregate endowment can be different in different states.

The first step in our analysis is a characterization of Pareto efficient allocations under external enforcement when all players are Bayesian expected utility maximizers. We say that the player i's prior is closer to the truth than player j's prior is if one of the following conditions is true: 1) player i learns a true data generating process faster than player j, 2) a true data generating process is possible under player i's prior, but is impossible under player j's prior, or 3) player i's asymptotic posterior belief is closer to a true data generating process then player j's asymptotic posterior belief is. If player i's prior is closer to the truth than player j's prior is, Schwartz (1965) and Berk (1965) imply that the ratio of player i' prior to player j's prior diverges almost surely. Pareto efficiency implies that if an event is more likely under player i's prior than it is under player j's prior, player i's consumption in that event should be higher than player j's consumption. Therefore, if player i's prior is closer to the truth than player j's prior is, Pareto efficiency means that player j's consumption converges to zero in the long run because the ratio of player i' prior to player j's prior diverges.

The second step in our analysis is a characterization of Pareto efficient subgame perfect equilibria of the risk sharing game when the players are Bayesian expected utility maximizers. We show that no player's consumption can fall below a certain threshold in a Pareto efficient subgame perfect equilibrium. The intuition behind this result is the following. Suppose that player i's current consumption is very small in a Pareto efficient subgame perfect equilibrium. Then, player i's self enforcement condition cannot be binding with probability one in the next period because otherwise she would have an incentive to deviate right now. Therefore, in some state s, player i's self enforcement condition is not binding. In this case, player i's consumption growth is limited.

We explain why player i's consumption growth is limited when her self enforcement condition is not binding. Choose player $j \neq i$. If player j's self enforcement condition is not binding in the state s, Pareto efficiency implies that player i' and j's consumptions in the state s are determined by their marginal rate of substitutions between two consecutive histories. Since player i's current consumption is small, her consumption in the state s is also small. If player j's self enforcement condition is binding in state s, player j's consumption is larger than what it would be if her self enforcement condition was not binding in state s. This is because if player j's consumption in state s had been determined by the marginal rate of substitutions, player j's self enforcement condition would have been violated. This means that when player j's self enforcement condition is binding in state s, player i's consumption would be lower than what it would be if player j's self enforcement condition was not binding in state s. This means that when player i's self enforcement condition is not binding, her consumption cannot increase by a large amount. This argument can continue if player i's current consumption is sufficiently small. In other words, once player i's consumption becomes very small, she expects low consumptions for long periods, and therefore, she has an incentive to deviate. Hence, every player's consumption is bounded from below in a Pareto efficient subgame perfect equilibrium.

The second step of our analysis implies that if player i's prior is closer to the truth than player j's prior is, no Pareto efficiency allocation under external enforcement is a Pareto efficient subgame perfect equilibrium of the risk sharing game even if players are sufficiently patient. If player i's prior is closer to the truth than player j's prior is, Pareto efficiency under external enforcement means that player j's consumption converges to zero. However, no player's consumption converges to zero in a Pareto efficient subgame perfect equilibrium.

The third step of the analysis is a characterization of Pareto efficient allocations when the players are ambiguity averse and have commitment power. We adopt the Gilboa-Schmeidler (1989) model for ambiguity aversion and adopt the recursive multiple priors model suggested by Epstein and Schneider (2003) to take care of dynamic consistency. That is, in our model, a player is ambiguity averse at every history. We assume that every player has a largest set of priors, which is the set of almost all stochastic processes. In this case, Pareto efficiency implies that each player could consume a constant amount of endowments in a state with a minimum in the state in which aggregate endowment is the lowest. In other words, if an allocation is Pareto efficient, each player's consumption depends only on a state, but is independent of periods. A player consumes the lowest amount when aggregate endowment is the lowest. An intuition for this result comes from certainty equivalent. For a given Pareto efficient allocation, one can find a certain equivalent for each player. In our model, sum of the players' certainty equivalents is less than or equal to present value of the lowest aggregate endowment. Clearly, the sum of certainty equivalents should be equal to the lowest endowment. So, when all players are ambiguity averse, one can find a payoff-equivalent stationary allocation, given a Pareto efficient allocation.

The final step of our analysis is a characterization of a subgame perfect equilibrium when the players are ambiguity averse. We focus on stationary allocations since we can find a stationary allocation given the Pareto efficient allocation under external enforcement. In our model every player's endowment depends only on the state. From this assumption, it is clear that if the players are sufficiently patient, a stationary allocation, which is payoff-equivalent to some Pareto efficient allocation, can be a subgame perfect equilibrium.

The rest of this paper is organized as follows. In Section 2, we illustrate an example to outline the basic structure of our model. In Section 3, we describe a risk sharing game. In Section 4, we analyze asymptotic behavior of Pareto efficient allocations under external enforcement, assuming that all players are expected utility maximizing Bayesian learners. In Section 5, we characterize a long run behavior of consumption in Pareto efficient subgame perfect equilibria when all players are Bayesian. In Section 6, we make some further comments on Pareto efficient subgame perfect equilibria with Bayesian players. In Section 7, we characterize Pareto efficient allocation when all players are ambiguity averse and have commitment power. In Section 8, we partially characterize a Pareto efficient subgame perfect equilibrium when all players are ambiguity averse. Section 9 is a conclusion.

2.1.1 Related Literature

Blume and Easley (2006) is closely related to this work. They consider an infinite horizon endowment economy in which traders have different priors on the true data generating process, and analyze the asymptotic behavior of Pareto efficient allocations in complete markets. As an example, Blume and Easley Section 3.1 (2006) analyzes independently and identically distributed economies in which traders do not learn over time.[2] They show that a trader's *survival index*, which is defined as a function of the trader's discount factor and the Kullback-Leibler divergence between the truth and the trader's prior, determines whether the trader survives in the long run. In our model, players can learn. Assuming that players have commitment power, we obtain a

[2]Blume and Easley (2006) refers to their unpublished manuscript for a complete discussion of i.i.d. economies. We contacted the authors, and could not have a chance to look at the manuscript.

similar result. We define "closeness" between the truth and a player's prior, and show that if player i's prior is "closer" to the truth than player j's prior is, player j does not survive in the long run.

In the comparison of the main theorems in Blume and Easley (2006) and our results, we would like to discuss two issues. The first issue is about absolute continuity. For instance, Blume and Easley Theorems 1 – 3 (2006) basically imply that a trader survives if and only if the true data generating process is absolutely continuous with respect to the trader's prior. However, our model covers cases in which the true data generating process is not absolutely continuous with respect to any player's prior, and talks about the long run survival. The second issue is about the analysis of the long run behavior of Pareto efficient allocations when the true data generating process is not included in the union of the support of every player's prior. Blume and Easley Theorems 1 – 6 apply when at least the support of one player's prior contains the true data generating process. In other words, the theorems are silent when the true data generating process is not included in the union of the support of all players' priors. However, in this paper, we analyze a subgame perfect equilibrium, and thus we need to specify what happens to every player's consumption even if the true data generating process is not an element of the union of the support of all players' priors. Focusing on i.i.d. processes with learning, we do not have to rely on absolute continuity and are able to describe the asymptotic behavior of a player's consumption, regardless of whether the true data generating process is included in the support of some player's prior.

Another stream of literatures on heterogeneous priors and market selection is focused on cases in which endowment follows a Geometric Brownian motion. Kogan *et al.*'s result (Proposition 2, 2006) states that an agent with a wrong belief never survives. Yan (2008) defines an agent's survival index, and Proposition 2 in the paper states that only those with the lowest survival index survive in the long run. Borovička (2008) uses a recursive utility and shows that under a certain specification of recursive preferences, an agent with an incorrect belief can survive. For those literatures, endowment is unbounded, so that the property of utility function is a determinant of the long run survival. Like Blume and Easley (2006), we assume that aggregate endowment is bounded, and as a result, the specification of utility function has no impact on the long run survival. Kogan *et al.* (2017) considers an economy in which there are two agents who have the same utility function but have different priors, and finds necessary conditions and sufficient conditions under which an agent survives in the long run. Since Kogan *et al.* (2017) covers cases in which endowment is unbounded, its results, such as Theorem 4.1 express conditions in terms of absolute risk aversion and the ratio of the two agents' priors.

Literature on model misspecification is also closely related to this paper. Schwartz (1965) and Berk (1965) study the limiting behavior of a Bayesian estimator when models are misspecified. They study whether a Bayesian estimator converges when the truth is not included in the support of a Bayesian learner's prior. It turns out that among all distributions in the support of the prior, the Bayesian learner asymptotically assigns probability one to the distribution that is closest to the truth in terms of Kullback-Leibler divergence. This result helps us analyze Pareto efficient allocations under external enforcement when Bayesian players learn in independently and identically distributed economies.

Kocherlakota (1996) studies properties of Pareto efficient subgame perfect equilibria of a risk sharing game in which two players have a common prior about an endowment process. Our risk sharing game is an extension of the Kocherlakota (1996) model in the sense that in our model, players have heterogeneous priors and learn over time. It is worth mentioning that the Kocherlakota (1996) arguments used to describe the dynamic paths of allocations in a Pareto efficient subgame perfect equilibrium are valid in our model. Applying similar arguments, we show that no player's consumption converges to zero in a Pareto efficient subgame perfect equilibrium.

Regarding ambiguity aversion, we use the maxmin criterion suggested in Gilboa and Schmeidler (1989). When an agent follows the maxmin criterion, she imagines the worst prior given the action. The expected utility is calculated using the given action and the worst prior.

Epstein and Schneider (2003) propose the recursive multiple priors model in order to handle an ambiguity averse agent in dynamic environments. They show that if an agent has a "rectangular" set of priors, dynamic consistency is obtained. The idea is that if the set of priors is rectangular, the agent can be thought of as Bayesian and the corresponding rule of updating the set of prior is Bayes' applied prior-by-prior. We adopt the recursive multiple priors model, and describe a class of Pareto efficient subgame perfect equilibrium when players are ambiguity averse. Epstein and Schneider (2007) study how an ambiguity averse decision maker can learn over time using their recursive multiple priors model. They introduce a way of reevaluating the set of priors. So, in their (2007) paper, the decision maker not only updates the set of priors but also selects priors through a "reevaluation process." Epstein and Schneider (2007) obtain a result similar to Schwartz (1965) and Berk (1965). That is, in their set-up, the ambiguity averse decision maker will believe that the truth is the one that is closest to the true data generating process among the possible beliefs.

Another stream of literature on ambiguity aversion is smooth ambiguity aversion, which was introduced in Klibanoff, Marinacci and Mukerji (2005). Klibanoff, Marinacci and Mukerji (2009) introduces recursive smooth ambiguity preferences to take care of dynamic consistency. The difference between the recursive multiple priors model and the recursive smooth ambiguity aversion is that an agent with the recursive multiple priors model does not learn the truth whereas an agent with the recursive smooth ambiguity can learn the truth. In this paper, we focus on the recursive multiple priors model.

There are several literatures on efficient allocation under ambiguity aversion, including Ghirardato and Siniscalchi (2016), Rigotti, Shannon and Strzalecki (2008), and Billot *et al.* (2000). In these papers, it is assumed that aggregate endowment is constant. The common result is that when all ambiguity averse agents share a common prior and markets are complete, Pareto efficiency implies full insurance, i.e., every agent consumes a constant amount of good in every state. In our model, we do not assume that aggregate endowment is constant. But, we assume that every player has a rectangular set of priors (unlike those literatures). This assumption is to minimize the impact of "closeness" between the truth and a player's prior. We show that if all players are ambiguity averse and have the largest rectangular set of priors, they can agree on a stationary allocation.

Regarding ambiguity aversion and market selection, there is a large volume of literature. Condie (2008) is closely related to this paper. He uses the recursive multiple priors model suggested by Epstein and Schneider (2003) to investigate whether ambiguity averse agents survive in the competitive market, assuming the existence of a Bayesian expected utility maximizer who knows the true data generating process. Condie Theorem 1 (2008) states that if there is aggregate uncertainty and there exists a Bayesian expected utility maximizer with the correct prior, any ambiguity averse agent who satisfies a certain condition vanishes. One difference between Condie (2008) and this paper is that we do not consider cases in which Bayesian expected utility maximizers and ambiguity averse players coexist. For Theorem 1 in Condie (2008), it is crucial to assume the existence of the Bayesian player who knows the truth. This is because an ambiguity averse agent with a rectangular set of priors can be thought of as Bayesian, and she has to compete against the Bayesian agent who knows the truth. As we mentioned in the previous paragraph, in our model, every ambiguity averse player can survive in the absence of Bayesian players. Da Silva (2011) studies the survival of an agent with variational preferences in the presence of an expected utility maximizing Bayesian agent with the correct beliefs, and Proposition 3 in the paper generalizes Condie's result (Theorem 1, 2008). Guerdjikova and Sciubba (2015) adopt the recursive smooth ambiguity aversion proposed by Klibanoff, Marinacci and Mukerji (2009), and analyze complete markets in which expected utility maximizing Bayesian learners and smooth ambiguity averse agents coexist. Guerdjikova and Sciubba's result (Proposition 5.12, 2015) states a sufficient condition under which a smooth ambiguity averse agent can survive even if there exists a Bayesian expected utility maximizer who knows the truth.

2.2 Example

We consider a simple risk sharing game. There are player 1 and player 2, who live for two periods. At the beginning of each period, Nature chooses a state from $S = \{A, B\}$. Each player believes that states are independently and identically distributed across periods. Player 1 and player 2 have different priors about states of world. Player 1 believes that state A is chosen with probability $q > 0.5$, and player 2 believes that state A is chosen with probability $1 - q$.

Each player receives one unit of a perishable good at the beginning of each period. The players have the same discount factor $\delta \in (0, 1)$ and the same utility function $u(c) = \ln c$ where c is consumption. In each period, after observing the state, each player can choose how much of her endowment to transfer to the other player in that period. Players make these decisions simultaneously.

If there was just one period, the unique Nash equilibrium is that both players make no transfer. Therefore, in the 2-period game, the unique subgame perfect equilibrium is to make no transfer in each period. However, the players can do better if they can agree on transfers, and if this agreement is enforced by a third party, such as a court. A good example is a complete market. This example can be thought of as a complete market when the players can commit. The following diagram shows the allocation in the complete market version of the example.

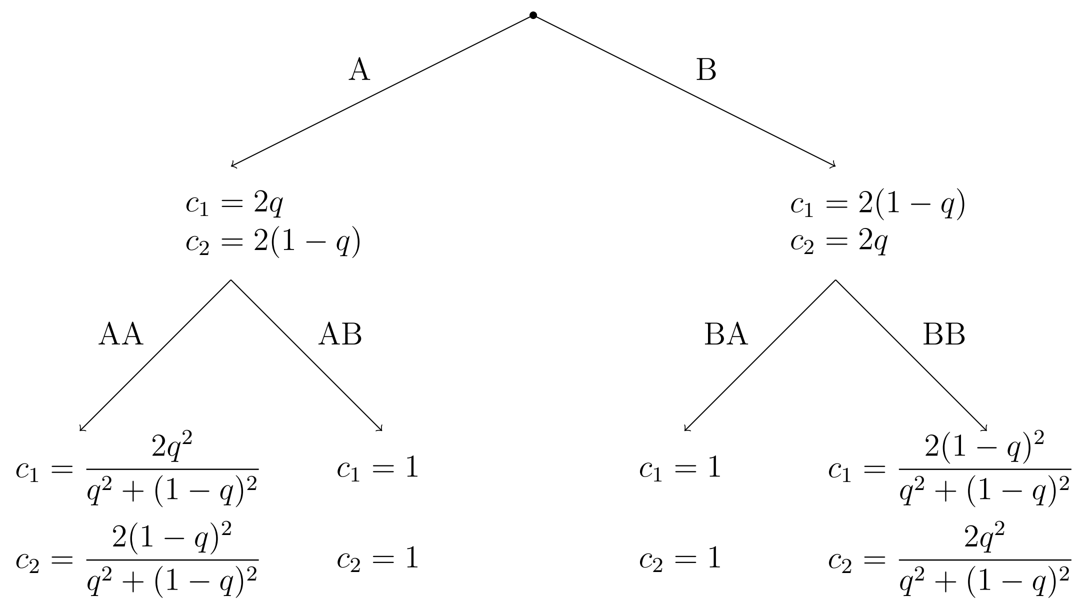

Player 2's expected utility is

$$\begin{aligned}
\mathbb{E}[u_2] &= (1-q)\ln 2(1-q) + q\ln 2q \\
&\quad + \delta\left[(1-q)^2 \ln\left(\frac{2(1-q)^2}{q^2+(1-q)^2}\right) + q^2 \ln\left(\frac{2q^2}{q^2+(1-q)^2}\right)\right] \\
&\quad > 0.
\end{aligned}$$

Since each player's expected utility is zero in the unique subgame perfect equilibrium, player 2 prefers the above allocation. Moreover, because the example is symmetric, the above allocation also gives player 1 an expected utility strictly larger than zero. Therefore, the two players would be willing to commit themselves to the above allocation. This argument means that existence of external enforcement has an impact on the players' agreement.

Ambiguity aversion is another factor that would have some impact on the players' agreement. If the two players behave according to maxmin criterion (Gilboa and Schmeidler (1989)), they would not agree on the above allocation. For simple argument, let us assume that every player thinks the probability of Nature choosing A can be any number in $[0, 1]$. Given the above allocation, the worst case from player 1's perspective is when Nature chooses B in every period. Since $q > 0.5$, if Nature always chooses B, player 1's consumptions in period 1 and period 2 will be lower than his endowment. Therefore, if player 1 calculates his expected utility according to maxmin criterion, he would not agree on the above allocation. Similarly, player 2 prefers her endowment to the above allocation.

To sum up, this example illustrates the idea that an allocation that the two players would agree on depends on whether external enforcement is available or not, and whether the two players are Bayesian or ambiguity averse. In this work, we generalize the example and study the impact of model uncertainty and commitment on outcomes that players would agree on.

2.3 Model: A Risk Sharing Game

Time is discrete, and there are infinitely many periods: $t = 1, 2, \ldots$. There are I infinitely-lived players who consume a single kind of a perishable and continuously divisible good. By abusing notation, I is also the set of players. In each period t, Nature chooses a state s_t from a finite set $S = \{1, \cdots, K\}$ of possible states. We postpone the description of the stochastic process that governs the selection of states until later in this section. Player i's endowment with the consumption good in period t is a function of the state $e_i : S \longrightarrow [\underline{e}, \overline{e}]$ for some $0 < \underline{e} < \overline{e}$. $e(s) = \sum_{i \in I} e_i(s)$ is the sum of the endowments in state s. Without loss of generality, we can assume that $e(1) \leq \cdots \leq e(K)$. Thus $e_i(s_t)$ is player i's endowment in period t, and $e(s_t)$ is the aggregate endowment in period t. The history of states up to period t is denoted by $s^t = (s_1, \cdots, s_t)$.

In each period t, after observing the state s_t, each player i can choose an amount $\tau_{ij}(s_t)$ of the good that he transfers to agent $j \neq i$. Player i can only choose non-negative transfers, and, moreover, his transfers have to satisfy:

$$\sum_{j \neq i} \tau_{ij}(s_t) \leq e_i(s_t).$$

The players simultaneously make transfers. Player i's consumption in period t is then:

$$c_i(s_t) = e_i(s_t) - \sum_{j \neq i} \tau_{ij}(s_t) + \sum_{j \neq i} \tau_{ji}(s_t).$$

Player i's utility from consumption in period t is given by $u_i(c(s_t))$, where $u_i : \mathbb{R}_+ \longrightarrow \mathbb{R}$ is player i's instantaneous utility function.[3] We assume that for all $i \in I$, u_i is monotonically increasing, strictly concave and continuously differentiable. In addition, we assume that for all $i \in I$, $\lim_{c \to 0+} u_i'(c) = \infty$. Given the infinite sequence of states $(s_1, s_2, \cdots)$, player i's overall utility is then the present discounted value of per period utilities:

$$\sum_{t=1}^{\infty} \delta^{t-1} u_i(c_i(s_t)).$$

Note that we assume that all players have the same discount factor. This is because differences in discount factors could lead to differences in players'

[3] $\mathbb{R}_+ = [0, \infty)$.

long-run survival, and we want to isolate the effect of differences in beliefs on long-run survival.

We now turn to a description of the stochastic process that determines the state in each period. $\mathcal{S}$ is the set of infinite sequences of states.

$$\mathcal{S} = S \times S \times S \times \cdots \equiv S^\infty.$$

An element $\omega \in \mathcal{S}$ is called a path. An element $s^t = (s_1, \cdots, s_t) \in S^t$ is called a state history. Clearly, S^t is the set of all state histories of length t. Given any state history $s^t \in S^t$, we construct a partition $\mathcal{F}_t$ of $\mathcal{S}$. For $s^t = (\bar{s}_1, \bar{s}_2, \ldots, \bar{s}_t) \in S^t$, we define

$$\mathcal{S}(s^t) = \{(s_1, s_2, s_3, \ldots) \in \mathcal{S} | s_\tau = \bar{s}_\tau \text{ for all } 1 \leq \tau \leq t\}.$$

The partition $\mathcal{F}_t$ is the collection of $\mathcal{S}(s^t)$. That is,

$$\mathcal{F}_t = \bigcup_{s^t \in S^t} \{\mathcal{S}(s^t)\}.$$

When $t = 0$, $\mathcal{F}_0 = \{\mathcal{S}\}$, and $s^0 \in \mathcal{F}_0$ is $\mathcal{S}$. Let $\mathcal{F}_\infty \equiv \cup_{t=1}^{\infty} \mathcal{F}_t$. $\mathcal{F}$ is the smallest sigma-algebra that contains $\mathcal{F}_\infty$. By abusing notation, s^t also refers to $\mathcal{S}(s^t)$.

The true data generating process is that states are i.i.d. Let ΔS be the set of all probability measures on S. For $q \in \Delta S$, ν_q is the induced measure on $(\mathcal{S}, \mathcal{F})$. Nature chooses a state s according to some $q \in \Delta S$ in every period. In other words, Nature chooses the path $\omega \in \mathcal{S}$ according to ν_q.

Players will also have beliefs about the data generating process. These beliefs will be probability measures on $(\mathcal{S}, \mathcal{F})$. Players will evaluate strategies by calculating the expected value of their discounted expected utility as defined above. Further details of players' evaluations of consumption plans will be specified in subsequent sections.

We now describe players' strategies. $\tau(s^t)$ denotes the transfer history, which is the tuple of all players' past transfers at the state history $s^t = (s_1, s_2, \cdots, s_t)$:

$$\tau(s^t) = \Big((\tau_{ij}(s_1)), \cdots, \tau_{ij}(s_{t-1}))\Big)_{i\in I, j\in I}$$

Player i's strategy σ_i specifies her transfers in every state history s^t as a function of the state history s^t and the transfer history $\tau(s^t)$. $\sigma_{ij}(s^t, \tau(s^t))$ denotes the transfer from player i to player j in period t. Σ_i is the set of player i's strategies, and $\Sigma = \Sigma_1 \times \cdots \times \Sigma_I$. An element $\sigma \in \Sigma$ is called a strategy profile.

Given the strategy profile σ, we can compute player i's transfers τ_{ij}^{σ} and consumption plan c_i^{σ}. For $s^1 \in \mathcal{F}_1$,

$$\begin{aligned}
\tau_{ij}^{\sigma}(s^1) &= \sigma_{ij}(s^0), \\
\tau^{\sigma}(s^1) &= \left(\tau_{ij}^{\sigma}(s^0)\right)_{i\in I, j\in I}.
\end{aligned}$$

For $s^2 = (s_1, s_2)$,

$$\begin{aligned}
\tau_{ij}^{\sigma}(s^2) &= \sigma_{ij}\left(s^2, \tau^{\sigma}(s^1)\right), \\
\tau^{\sigma}(s^2) &= \left(\tau_{ij}^{\sigma}(s^1), \tau_{ij}^{\sigma}(s^2)\right)_{i\in I, j\in I}.
\end{aligned}$$

By induction, we can compute $\tau_{ij}^{\sigma}(s^t)$ for all $s^t \in \mathcal{F}_{\infty}$.

Using the transfers τ_{ij}^{σ}, we can compute player i's consumption in every state history $c_i^{\sigma}(s^t)$. For computation of c_i^{σ}, we assume that every player follows σ. If external enforcement is available, this assumption is automatically satisfied. In the absence of external enforcement, the assumption would be meaningful if σ is a subgame perfect equilibrium, and c_i^{σ} will be player i's consumption along the equilibrium path. Since no player has an incentive to discard the good, for every $s^t \in \mathcal{F}_t$

$$c_i^\sigma(s^t) = e_i(s^t) - \sum_{j \neq i} \tau_{ij}^\sigma(s^t) + \sum_{j \neq i} \tau_{ji}^\sigma(s^t).$$

Player i's utility given strategy profile σ, and given a sequence of states, is then:

$$\sum_{t=1}^{\infty} \delta^{t-1} u_i(c_i^\sigma(s^t)).$$

2.4 Bayesian Players and External Enforcement

In this section, we analyze the asymptotic behavior of Pareto efficient allocations when all players are Bayesian and external enforcement is available. By external enforcement, we mean that given the strategy profile σ, the transfers $\tau_{ij}^\sigma(s^t)$ are enforced by an external enforcer, such as a court, in every state history. In other words, external enforcement means that player i must consume $c_i^\sigma(s^t)$ in every state history s^t.

The main result in this section is that if player i's prior is *closer* to the truth than player j's prior is, player j's consumption converges to zero. To prove this result, we first define a Bayesian player and Pareto efficiency under external enforcement. Then, we introduce a notion of "closeness" between the truth and a player's prior. Applying the limiting behavior of Bayes estimators, we can establish the main result.

We start with the definition of a Bayesian player.

Definition 2.4.1. *Player i is said to be Bayesian if she has a prior belief*

μ_i, a probability measure defined on $(\mathcal{S}, \mathcal{F})$, that is of the form:

$$\mu_i = \int\limits_{\Delta S} \nu_q F_i(dq).$$

for some $F_i \in \Delta(\Delta S)$ and evaluates consumption plans according to her expected utility.

Definition 2.4.1 reflects that players believe that the process is i.i.d. Therefore, their prior is assumed to be the convex combination of i.i.d processes. For example, suppose that $S = \{A, B\}$. Let q be the probability of state A, and assume that player i knows that $q = 0.3$ and $q = 0.6$ are equally likely. In this case, $F_i(0.3) = F_i(0.6) = 0.5$ and F_i is zero otherwise. Player i's prior is then:

$$\mu_i = 0.5\nu_{(0.3)} + 0.5\nu_{(0.6)}.$$

For technical convenience, we assume that for each $i \in I$, the support of F_i is a compact subset of the interior of ΔS. This means that no player believes that the i.i.d process might be such that some state occurs with zero probability. We also assume that, if for an open set $W \subset \text{supp}(F_i)$, F_i conditional on W admits a density $f_i|_W$ with respect to Lebesgue measure, $f_i|_W(q) > 0$ for all $q \in W$.

Given the consumption plan c_i and the prior μ_i, we use V_i to denote player i's expected utility:

$$V_i(c_i, \mu_i) = \sum_{t=1}^{\infty} \sum_{s^t \in \mathcal{F}_t} \delta^{t-1} u_i(c_i(s^t))\mu_i(s^t).$$

Now, we turn to define Pareto efficiency under external enforcement.

Definition 2.4.2. *A strategy profile σ is Pareto efficient under external en-*

forcement if there is no strategy profile σ' such that for all $i \in I$,

$$V_i(c_i^{\sigma}, \mu_i) \leq V_i(c_i^{\sigma'}, \mu_i)$$

and for some $i \in I$, the above relationship holds with strict inequality.

In order to state our results, we need to introduce a notion of closeness between a true data generating process and a player's prior. We use the following notations. For $\omega \in \mathcal{S}$ and $t \in \mathbb{N}$ define:

$$\hat{q}_t(\omega) = \frac{1}{t}\left(n_{1,t}(\omega), \cdots, n_{K,t}(\omega)\right) \equiv \left(\hat{q}_{1,t}(\omega), \cdots, \hat{q}_{K,t}(\omega)\right)$$

where $n_{s,t}(\omega)$ is the number of realizations in the first t realizations in ω that equal state s. For $q \in \Delta S$, define

$$\mathcal{S}(q) = \left\{\omega \in \mathcal{S} \middle| \lim_{t\to\infty} \hat{q}_t(\omega) = q\right\}.$$

This is the set of infinite sequences where the frequency $\hat{q}$ converges to q.

For $q, q' \in \Delta S$, $D(q, q')$ the Kullback-Leibler divergence is defined as follows (Kullback p. 6, 1959):

$$D(q, q') = \sum_{s\in S} q(s)\left[\ln q(s) - \ln q'(s)\right].$$

Using $\mathcal{S}(q)$ and Kullback-Leibler divergence, we define a comparison of the true data generating process and two players' priors as follows.

Definition 2.4.3. *For $i \neq j$, μ_i is closer to $q \in \Delta S$ than μ_j is if:*

1. $\mu_i(\mathcal{S}(q)) > 0$ *and* $\mu_j(\mathcal{S}(q)) = 0$, *or*
2. $\min_{q'\in supp(F_i)} D(q, q') < \min_{q'\in supp(F_j)} D(q, q')$.

Let us interpret Definition 2.4.3 and discuss the implication of the conditions in the definition for players' learning. If the first condition holds, player i assigns positive probability to $\mathcal{S}(q)$ whereas player j assigns zero probability to $\mathcal{S}(q)$. In this case, it is certain that player i's posterior belief converges to ν_q. That is, player i learns the truth in the long run. There are two cases in which $\mu_j(\mathcal{S}(q)) = 0$: 1) q is not in the support of player j's prior, or 2) player j's prior has some positive density at q, but q is not an atom of player j's prior. If player j believes that q is impossible, she does not learn the truth in the long run, whereas player i learns the truth. If player j has some positive density at q, she learns the truth in the long run. However, we will show below in Lemma 2.4.1 that player i's posterior belief converges to q faster than player j's posterior belief does. In words, player i learns the truth faster than player j does.

Let us look at the second condition. We consider two cases. The first case is when $\min_{q' \in \text{supp}(F_i)} D(q, q') = 0$ and $\min_{q' \in \text{supp}(F_j)} D(q, q') > 0$. In this case, the true data generating process is included in the support of player i's prior, but the true data generating process is not in the support of player j's prior. This means that in the long run, player i learns the true data generating process whereas player j does not.

The second case is when $0 < \min_{q' \in \text{supp}(F_i)} D(q, q') < \min_{q' \in \text{supp}(F_j)} D(q, q')$. In this case, the true data generating process is in neither the support of player i's prior nor the support of player j's prior. This implies that in the long run, neither player i nor player j learns the truth. Schwartz (Theorem 6.1, 1965) and Berk (Theorem, p. 54, 1965) establish that the posterior belief converges to one that is closest to the true data generating process among possible beliefs in terms of Kullback-Leibler divergence. Therefore, in this case, the Kullback-Leibler divergence between player i's asymptotic posterior belief and the truth is strictly smaller than the Kullback-Leibler divergence between player j's asymptotic posterior belief and the truth.

The following lemma will be useful in characterizing Pareto efficient allocation under external enforcement.

Lemma 2.4.1. *Suppose that μ_i is closer to $q \in \Delta S$ than μ_j is. Then, $\frac{\mu_i}{\mu_j}$ goes to infinity ν_q-almost surely.*

$$\nu_q\left(\left\{(s_1, s_2, \cdots) \in \mathcal{S}(q) \middle| \lim_{t\to\infty} \frac{\mu_i(s^t)}{\mu_j(s^t)} = \infty\right\}\right) = 1.$$

(proof) See the Appendix.

We now write down the expected utility maximization problem that a Pareto efficient allocation must solve.

$$\max_{c_2,\ldots,c_I} \sum_{t=1}^{\infty} \sum_{s^t \in \mathcal{F}_t} \delta^{t-1} u_1\left(e(s^t) - \sum_{i=2}^{I} c_i(s^t)\right) \mu_1(s^t)$$

subject to

$$\sum_{t=1}^{\infty} \sum_{s^t \in \mathcal{F}_t} \delta^{t-1} u_i(c_i(s^t))\mu_i(s^t) \geq v_i \quad \forall i = 2, \cdots, I.$$

The first order condition implies that there exists $\lambda_i > 0$ for each $i \in I$ such that for all s^t,

$$\frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))} \frac{\mu_i(s^t)}{\mu_j(s^t)} = \frac{\lambda_i}{\lambda_j} \text{ for all } i, j \in I. \tag{2.1}$$

One way of understanding the first order condition is to think of different state histories as different "states." $\mathcal{F}_\infty$ can be thought of as a "new set of states." For player i's prior μ_i, one can generate a "new measure" $\tilde{\mu}_i$ in the following way. For every element $s^t \in \mathcal{F}_\infty$,

$$\tilde{\mu}_i(s^t) = (1-\delta)\delta^t \mu_i(s^t).$$

Note that $\sum_{s^t \in \mathcal{F}_\infty} \tilde{\mu}_i(s^t) = 1$. Pareto efficiency implies that for every pair of "new states" $s^t, s'^{t'}$ and for every pair of i, j, player i's marginal rate of substitution between s^t and $s'^{t'}$ is equal to player j's marginal rate of substitution between s^t and $s'^{t'}$. That is,

$$\frac{u_i'(c_i(s^t))\tilde{\mu}_i(s^t)}{u_i'(c_i(s'^{t'}))\tilde{\mu}_i(s'^{t'})} = \frac{u_j'(c_j(s^t))\tilde{\mu}_j(s^t)}{u_j'(c_j(s'^{t'}))\tilde{\mu}_j(s'^{t'})}.$$

This means that for every i, there exist $\lambda_i > 0$ such that for every $s^t \in \mathcal{F}_\infty$,

$$\frac{u_i'(c_i(s^t))\tilde{\mu}_i(s^t)}{u_j'(c_j(s^t))\tilde{\mu}_j(s^t)} = \frac{\lambda_i}{\lambda_j}.$$

Clearly, the above equation is the same as the first order condition (2.1). The new equation implies that the ratio of player i and j's probability weighted marginal utilities from consumption needs to be the same in all states. If this were not the case, then a Pareto improvement could be achieved by increasing player i's consumption and decreasing player j's consumption in the state in which the ratio is higher, and changing consumption in the opposite direction in the state in which the ratio is lower.

Another way of understanding the first order condition (2.1) is to consider the dynamics of the consumptions between two consecutive state histories. Suppose that the players arrived at the state history s^t, but they do not consume yet. Choose two players i and j. They are about to consume $c_i(s^t)$ and $c_j(s^t)$, respectively. Pareto efficiency means that player i and player j have no incentive to trade between two consecutive periods. In other words, Pareto efficiency requires that player i''s marginal rate of substitution between s^t and (s^t, s) be equal to player j''s marginal rate of substitution between s^t and (s^t, s). This means that for every $s \in S$,

$$\frac{u_i'(c_i(s^t))}{u_i'(c_i(s^t, s))\mu_i(s^t, s|s^t)} = \frac{u_j'(c_j(s^t))}{u_j'(c_j(s^t, s))\mu_j(s^t, s|s^t)}.$$

After rearranging, we obtain

$$\frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))} = \frac{u_i'(c_i(s^t, s))\mu_i(s^t, s|s^t)}{u_j'(c_j(s^t, s))\mu_j(s^t, s|s^t)}. \tag{2.2}$$

Multiplying both sides by $\frac{\mu_i(s^t)}{\mu_j(s^t)}$ results in the first order condition (2.1).

So far, we have explained the first order condition (2.1) relying on the equalization of marginal rates of substitution that is familiar as a necessary condition for Pareto efficiency in general equilibrium theory. However, we will consider the dynamic path of allocations later, and it would be more useful to use the equation (2.2) because the left hand side depends only on s^t and the right hand side depends only on (s^t, s). We will constantly use the ratio of two players' marginal utilities to describe the dynamics of consumptions. As an exercise, let us explain the implication of the equation (2.2). Suppose that the equation (2.2) does not hold and that the left hand side is larger than the right hand side. In this case, we can find a Pareto improvement. Let us increase $c_i(s^t)$ by ϵ_1 and decrease $c_i(s^t, s)$ by ϵ_2 such that player i receives the same expected utility up to the first order. That is, ϵ_1 and ϵ_2 satisfy the following equation.

$$u_i'(c_i(s^t))\epsilon_1 - \delta u_i'(c_i(s^t, s))\mu_i(s^t, s|s^t)\epsilon_2 = 0.$$

Now, let us look at the change in player j's expected utility. Her consumption in the state history decreases by ϵ_1 and her consumption in (s^t, s) increases by ϵ_2.

$$
\begin{aligned}
&- u_j'(c_j(s^t))\epsilon_1 + \delta u_j'(c_j(s^t,s))\mu_j(s^t,s|s^t)\epsilon_2 \\
&= u_j'(c_j(s^t))\epsilon_1 \left[-1 + \frac{\delta u_j'(c_j(s^t,s))\mu_j(s^t,s|s^t)}{u_j'(c_j(s^t))}\frac{\epsilon_2}{\epsilon_1}\right] \\
&= u_j'(c_j(s^t))\epsilon_1 \left[-1 + \frac{\delta u_j'(c_j(s^t,s))\mu_j(s^t,s|s^t)}{u_j'(c_j(s^t))}\frac{u_i'(c_i(s^t))}{\delta u_i'(c_i(s^t,s))\mu_i(s^t,s|s^t)}\right] \\
&= u_j'(c_j(s^t))\epsilon_1 \left[-1 + \frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))}\frac{u_j'(c_j(s^t,s))\mu_j(s^t,s|s^t)}{u_i'(c_i(s^t,s))\mu_i(s^t,s|s^t)}\right]
\end{aligned}
$$

Therefore, if the ratio of two players' marginal utility from today consumption is larger than the ratio of their probability-weighted-marginal utilities from tomorrow consumption, the above quantity is positive. This means that one can find a Pareto improvement if the ratio of two players' marginal utility from today consumption is larger than the ratio of their probability-weighted-marginal utility from tomorrow consumption. Therefore, if an allocation is Pareto efficient under external enforcement, for every pair of players, the ratio of their marginal utilities from today consumption is equal to the ratio of their probability-weighted-marginal utilities from tomorrow consumption.

From the first order condition (2.1), we can see that if $\frac{\mu_i}{\mu_j}$ goes to infinity, $\frac{u_i'}{u_j'}$ converges to zero. That is, player j's consumption converges to zero when $\frac{\mu_i}{\mu_j}$ goes to infinity. This idea is reflected in the following proposition.

Proposition 2.4.1. *Suppose that a strategy profile σ is Pareto efficient under external enforcement. If μ_i is closer to q than μ_j is, c_j converges to zero ν_q-almost surely.*

(proof) See the Appendix.

The intuition for Proposition 2.4.1 comes from the combination of Definition 2.4.3 and Lemma 2.4.1. Suppose $q \in \Delta S$ is the true data generating process. Proposition 2.4.1 implies that player j's asymptotic consumption converges to zero with probability 1 in the following cases: 1) player i learns the truth faster than player j, 2) the true data generating process is possible under player i's prior whereas it is impossible under player j's prior, and 3) neither player i nor player j learns the truth, but player i's asymptotic posterior belief is closer to the true data generating process than player j's asymptotic posterior belief is. We discuss these cases in the next paragraph.

The first case in the previous paragraph means that learning speed matters in terms of long run survival. Proposition 2.4.1 implies that in our model, if player i learns the truth faster than player j, player j's consumption gets smaller. The second case means that if player i learns the truth, player j needs to learn the truth in order to survive in the long run. It does not matter whether player i survives in the long run. The third case means that if no player learns the truth, the survivors are the players who have the prior "closest" to the truth. In other words, it is important for a player to have the least incorrect belief in order to survive in the long run. It would be worth mentioning that the third case means that a player being a winner in the long run does not necessarily mean that she knows the truth.

We illustrate the third case using the Example in Section 2. The example needs to be modified to deal with the infinite horizon. Suppose that Nature always chooses state A. For player 1, the probability of observing t A's in a row is q^t. From player 2's perspective, the probability of observing t A's in a row is $(1-q)^t$. The ratio of those two probabilities is $(\frac{q}{1-q})^t$, which diverges to infinity as t goes to infinity. Note that neither of the two players learns the truth because they have point priors. However, Pareto efficiency implies that player 2's asymptotic consumption is very small. We can show this using Kullback-Leibler divergence. For player 1, the Kullback-Leibler divergence between the truth and her prior is $-\ln q$. For player 2, the Kullback-Leibler divergence between the truth and his prior is $-\ln(1-q)$. Since $q > 0.5 > 1-q$, $-\ln q < -\ln(1-q)$, this means that player 1's prior is closer to the truth than player 2's prior.

Finally, we would like to compare Proposition 2.4.1 and the Blume and Easley results in Section 3.1 (2006). One difference between Blume and Easley Section 3.1 and our model is that in Blume and Easley (2006), traders have different discount factors. We can relax the assumption that every player has same discount factor, and obtain the result similar to Blume and Easley result. Let δ_i be player i's discount factor. The first order condition (2.1) becomes

$$\frac{\delta_i^{t-1}}{\delta_j^{t-1}} \frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))} \frac{\mu_i(s^t)}{\mu_j(s^t)} = \frac{\lambda_i}{\lambda_j}.$$

So, long run survival depends on the ratio of player i's discount factor to player j's discount factor as well.

Another difference between Blume and Easley (2006) Section 3.1 and our model is that traders in Blume and Easley (2006) do not learn, whereas in our model players can learn. So, our model can be thought of as a generalization of their i.i.d. example.

Blume and Easley Theorems 1 – 3 (2006) rely on the notion of absolutely continuity. However, our model covers cases in which the true data generating process is not absolutely continuous with respect to some player's prior, and talks about the long run survival.

Blume and Easley Theorems 1 – 6 (2006) apply when at least the support of one player's prior contains the true data generating process. In other words, the theorems do not tell much about long run survival when the true data generating process is not included in the union of the support of all players' priors. However, in this paper, we analyze a subgame perfect equilibrium, and thus we need to specify what happens to every player's consumption even if the true data generating process is not included in the union of the support of all players' priors. Focusing on i.i.d. processes with learning, we do not rely on absolute continuity. Also, we are able to describe the asymptotic behavior of a player's consumption, regardless of whether the true data generating process is included in the support of some player's prior.

2.5 Bayesian Players and Self Enforcement

In this section, we analyze a property of Pareto efficient subgame perfect equilibria of the risk sharing game when the players are Bayesian. Our main result in this section is that for each player, there exists a lower bound such that her consumption does not fall below the lower bound in every Pareto efficient subgame perfect equilibrium. Roughly speaking, no player starves in the long run if the players implement a Pareto efficient subgame perfect equilibrium. This implies that even if player i's prior is closer to the true data generating process than player j's prior is, player j's consumption does not converge to zero. One implication of the main result is that no Pareto efficient strategy profile under external enforcement is a Pareto efficient subgame perfect equilibrium. In this sense, the value of external enforcement is positive.

To derive the main result, we first define a subgame perfect equilibrium, and completely characterize subgame perfect equilibriums. The characterization of subgame perfect equilibrium is the following. The worst subgame perfect equilibrium is that every player makes zero transfer to everyone else, regardless of how other players have played. This is because the unique Nash equilibrium of the stage game is that every player makes zero transfer, and every player receives their minimax payoffs in the unique Nash equilibrium. This property allows us to use a grim trigger strategy to characterize a subgame perfect equilibrium.

After the characterization of all subgame perfect equilibria, we define Pareto efficiency without commitment. To understand properties of the dynamic paths of consumption, we set up a standard expected utility maximization problem with constraints and derive the first order condition. The main result in this section is derived from the first order condition.

We begin with defining a subgame perfect equilibrium.

Definition 2.5.1. *A strategy profile σ is a subgame perfect equilibrium of the risk sharing game if it induces a Nash equilibrium of every subgame.*

To characterize subgame-perfect equilibria we need to consider a player's continuation payoff in the state history s^t. Given the strategy profile σ and the prior μ_i, player i's continuation payoff at the state history s^t is denoted by $V_i(c_i^\sigma, \mu_i | s^t)$:

$$V_i(c_i^\sigma, \mu_i | s^t) = \sum_{r=1}^{\infty} \delta^{r-1} u_i(c_i^\sigma(s^{t+r}))\mu_i(s^{t+r}|s^t).$$

We turn to characterize subgame perfect equilibria. Note that the static risk sharing game, that is, the risk sharing game with just one period, is similar to the Prisoner's dilemma. In the static risk sharing game the players can expect higher utilities if before observing the state, they commit themselves to state contingent transfers. However, the unique Nash equilibrium of the static risk sharing game is that every player makes zero transfer. This observation is useful in characterizing a subgame perfect equilibrium.

$\underline{\sigma}$ denotes the "no-transfer" strategy profile under which for every i, j, s^t and $\tau(s^t)$, $\tau_{ij}^{\underline{\sigma}}(s^t, \tau(s^t)) = 0$. Clearly, this means that for every $i \in I$ and every state history $s^t \in \mathcal{F}_\infty$, $c_i^{\underline{\sigma}}(s^t) = e_i(s^t)$.

Lemma 2.5.1. *The "no-transfer" strategy profile $\underline{\sigma}$ is a subgame perfect equilibrium. And, in any subgame perfect equilibrium σ, $V_i(c_i^\sigma, \mu_i | s^t) \geq V_i(e_i, \mu_i | s^t)$ for every $i \in I$ and every $s^t \in \mathcal{F}_\infty$.*

Proof. In the static risk sharing game, the unique Nash equilibrium is that the I players make zero transfer. Therefore, it is a subgame perfect equilibrium for the players to make zero transfer regardless of how the players have played in the past. This argument proves that the no-transfer strategy profile is a subgame perfect equilibrium.

To prove the second statement, we calculate minimax payoffs first. The harshest way to punish player i is that all other players make zero transfer to player i. For player i, the best response to receiving zero transfer is to make zero transfer to all other players. This means that for each player, consuming their own endowments is minimax payoff, which in turn implies that no-transfer strategy profile is the worst subgame perfect equilibrium. □

Using Lemma 2.5.1, we can use a grim trigger strategy profile to construct all subgame perfect equilibria.

Lemma 2.5.2. *For an allocation $(c_1, \cdots, c_I)$, there exists a subgame perfect equilibrium σ with $c_i^\sigma = c_i$ for all i if and only if for all $i \in I$ and for $s^t \in \mathcal{F}_\infty$,*

$$u_i(c_i(s^t)) + \delta V_i(c_i, \mu_i | s^t) \geq u_i(e_i(s^t)) + \delta V_i(e_i, \mu_i | s^t).$$

(proof) See the Appendix.

Our interest in this section is in Pareto efficient subgame perfect equilibria. To clarify the meaning of Pareto efficiency in the context of subgame perfect equilibria, we formally define a Pareto efficient subgame perfect equilibrium first.

Definition 2.5.2. *A strategy profile σ is a Pareto efficient subgame perfect equilibrium if it is a subgame perfect equilibrium and there is no subgame perfect equilibrium σ' such that for all $i \in I$,*

$$V_i(c_i^\sigma, \mu_i) \leq V_i(c_i^{\sigma'}, \mu_i)$$

and for some $i \in I$, the above relationship holds with strict inequality.

Note that Definition 2.5.2 is ex-ante Pareto efficiency.

To analyze Pareto efficient subgame perfect equilibria, we are going to set up an expected utility maximization problem. The problem we need to solve is the following:

$$\max_{c_2,\cdots,c_I} \sum_{t=1}^{\infty} \sum_{s^t \in \mathcal{F}_t} \delta^{t-1} u_1 \left(e(s^t) - \sum_{i=2}^{I} c_i(s^t) \right) \mu_1(s^t)$$

subject to

$$\sum_{t=1}^{\infty} \sum_{s^t \in \mathcal{F}_t} \delta^{t-1} u_i(c_i(s^t)) \mu_i(s^t) \geq v_i, \quad \forall i = 2, \cdots, I,$$

and for all $i = 1, 2, \cdots, I$ and all s^t

$$u_i(c_i(s^t)) + \delta V_i(c_i, \mu_i | s^t) \geq u_i(e_i(s^t)) + \delta V_i(e_i, \mu_i | s^t) \qquad \text{SEC}(i)$$

where SEC(i) is player i's self enforcement condition.

Using the first order condition equation, we can briefly describe the dynamics of players' consumptions. Choose the state history $s^t = (s^{t-1}, s)$. If player i and j's self enforcement conditions are not binding in s^t, then first order condition means that

$$\frac{u_i'(c_i(s^{t-1}))}{u_j'(c_j(s^{t-1}))} \frac{\mu_i(s^{t-1})}{\mu_j(s^{t-1})} = \frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))} \frac{\mu_i(s^t)}{\mu_j(s^t)}.$$

Equivalently,

$$\frac{u_i'(c_i(s^{t-1}))}{u_j'(c_j(s^{t-1}))} = \frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))} \frac{\mu_i(s^t | s^{t-1})}{\mu_j(s^t | s^{t-1})}.$$

This means that if the two players' self enforcement conditions are not binding, then the marginal rate of substitution between s^{t-1} and s^t should be the same. Note that this is equivalent to the first order condition (2.2).

If player i's self enforcement condition is binding and player j's self enforcement condition is not binding in s^t,

$$\frac{u_i'(c_i(s^{t-1}))}{u_j'(c_j(s^{t-1}))} > \frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))}\frac{\mu_i(s^t|s^{t-1})}{\mu_j(s^t|s^{t-1})}.$$

If player i's self enforcement condition is binding and player j's self enforcement condition is not binding in s^t, that means that $c_i(s^t)$ is higher than what it would be if the two self enforcement conditions are not binding in s^t. This means that $\frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))}$ is lower compared to the case in which SEC(i) and SEC(j) are not binding in s^t.

Both cases imply that if player j's self enforcement condition is not binding in s^t, player j's consumption increase from s^{t-1} to s^t is limited. Especially, if $c_j^\sigma(s^{t-1})$ is very small, and if player j's self enforcement condition is not binding in the state history (s^{t-1}, s), $c_j^\sigma(s^{t-1}, s)$ is also very small. This argument is used to prove the following proposition.

Proposition 2.5.1. *Fix $\delta < 1$. For each $i \in I$, there exists $\underline{c_i} > 0$ such that for every state history s^t and for every Pareto efficient subgame perfect equilibrium σ, $c_i^\sigma(s^t) \geq \underline{c_i}$.*

(proof) See the Appendix.

Proposition 2.5.1 implies that every player's consumption is bounded from below in a Pareto efficient subgame perfect equilibrium. Proposition 2.5.1 also means that a lower bound for a player's consumption is independent of Pareto efficient subgame perfect equilibria. In other words, a player's minimum consumption is guaranteed regardless of Pareto efficient subgame perfect equilibria.

We discuss how a lower bound for a player's consumption behaves as the discount factor δ goes to one. Following the proof of Proposition 2.5.1, one can construct a lower bound given the discount factor. This lower bound converges to zero as the discount factor goes to one. However, the proof of Proposition 2.5.1 does not show that the lower bound is tight. It is an open question to prove or disprove that a tight lower bound converges to zero as the discount factor goes to one.

Let us discuss the idea behind Proposition 2.5.1. Suppose not. Then, we can choose a Pareto efficient subgame perfect equilibrium σ in which $c_j^\sigma(s^t)$ is close to zero for some $j \in I$ and some state history s^t. Since everyone's consumption cannot be close to zero at the same time, there must exist player i whose consumption is large in s^t. Note that SEC(j) cannot be binding in period $t+1$ with probability one because otherwise SEC(j) is violated in s^t. Therefore, SEC(j) is not binding in some state s_{t+1} in period $t+1$. The first order condition implies that player j's consumption in $s^{t+1} = (s^t, s_{t+1})$ is very small because her consumption in s^t is very small. Now, we can repeat this argument. If the state s_{t+1} is realized in period $t+1$, player j's consumption is very small, which means that SEC(j) cannot be binding with probability one in period $t+2$, which in turn means that if SEC(j) is not binding in state s_{t+2}, player j's consumption in $s^{t+2} = (s^{t+1}, s_{t+2})$ is low . This logic means that if $c_j^\sigma(s^t)$ is sufficiently small, player j expects low consumption to be persistent in the future when she arrives at the state history s^t. However, this would mean that player j's continuation payoff in s^t is lower than her minimax payoff. Therefore, it would be impossible to choose a Pareto efficient subgame perfect equilibrium in which some player's consumption can be arbitrarily close to zero.

Proposition 2.5.1 has an implication on Pareto efficient allocations under external enforcement. We know from Proposition 2.4.1 that if player i's prior is closer to the truth than player j's prior is, there exists a path $\omega \in \mathcal{S}$ along which player j's consumption converges to zero. However, Proposition 2.5.1 means that in a Pareto efficient subgame perfect equilibrium, every player's consumption is bounded from below in every state history. This implies that if a strategy profile σ is Pareto efficient under external enforcement, it is not a Pareto efficient subgame perfect equilibrium. This idea is reflected in the following corollary.

Corollary 2.5.1. *Suppose 1) that all I players are Bayesian, 2) that μ_i is closer to q than μ_j is for some i, j and $q \in \Delta S$, and 3) that $\delta < 1$. Then, no Pareto efficient strategy profile under external enforcement is a Pareto efficient subgame perfect equilibrium.*

Proposition 2.5.1 and Corollary 2.5.1 hold as long as δ is strictly less than one. There are two interesting questions that I have not been able to address yet. The first question is whether Proposition 2.5.1 and Corollary 2.5.1 are true when $\delta = 1$, and the second question is whether the Pareto efficient frontier of the set of payoff vectors in subgame perfect equilibria converges to the Pareto efficient frontier of the set of payoff vectors with commitment.

Nonetheless, we can discuss Proposition 2.5.1 and Corollary 2.5.1. If the answer to the first question is yes, then the answer to the second question is no. No player starves in the long run even if the discount factor is one. External enforcement is valuable. If the answer to the first question is no, some player would starve in the long run. The set of payoff vectors of Pareto efficient subgame perfect equilibrium would converge to the set of payoff vectors of Pareto efficient allocations with commitment. However, if every player does not discount future payoffs at all, the analysis would not be very interesting.

For some class of repeated games such as some repeated Prisoner's dilemma games, there exists $\underline{\delta}$ such that for all $\delta \geq \underline{\delta}$, a Pareto efficient payoff vector with commitment can be achieved in a subgame perfect equilibrium. Proposition 2.5.1 and Corollary 2.5.1 imply that in our model, there is no such $\underline{\delta}$. So long as $\delta < 1$, the Pareto efficient frontier of the set of payoff vectors in subgame perfect equilibria is never equal to the frontier of the set of Pareto efficient payoff vectors with commitment. To sum up, as long as players are impatient, no player vanishes and the value of external enforcement is positive.

2.6 Further Comments on Bayesian Players and Self Enforcement

Proposition 2.5.1 is meaningful only if we know that a Pareto efficient subgame perfect equilibria exist. But this is easily seen, as the following argument shows. First, suppose we know that the set of subgame-perfect equilibrium payoffs is non-empty and compact. Then we could maximize the sum of all players' payoffs on this set, and the subgame-perfect equilibrium that corresponds to the argmax would, by construction, be Pareto efficient. Therefore, it is sufficient to argue that the set of subgame-perfect equilibrium payoffs is non-empty and compact. To see that it is non-empty, we have to give an example of a subgame-perfect equilibrium. The strategy profile in which no player makes any transfers at all, regardless of history, is such an subgame-perfect equilibrium. The subgame-perfect equilibrium payoffs are bounded from below by the payoffs from the no-transfer equilibrium. They are bounded from above because aggregate endowment is bounded. Finally, the closeness of the set of subgame-perfect equilibrium payoffs follows by adapting standard arguments that show the closeness of the set of subgame-perfect equilibrium payoffs in deterministic repeated games with perfect monitoring (Proposition 2.5.2 in Mailath and Samuelson (2006)).

One might be interested in further properties of the consumption dynamics in Pareto efficient subgame perfect equilibria, beyond what is described in Proposition 2.5.1. This is somewhat orthogonal to the main purpose of this chapter, and requires moreover either a detailed numerical study, or a somewhat complicated mathematical study. Therefore, we have not pursued this topic further.

2.7 Ambiguity Averse Players and External Enforcement

In this section, we analyze Pareto efficient allocations when all I players are ambiguity averse. The main result of this section is that if every player is ambiguity averse, they could agree on a stationary allocation and therefore, no one would starve in the long run.

By an ambiguity averse player, we mean that the player has a set of priors and chooses the worst prior given the consumption plan. This idea is originally suggested by Gilboa and Schmeidler (1989). In order to model a belief-updating-system and address the issue of dynamic consistency, we adopt the recursive-multiple-priors model proposed by Epstein and Schneider (2003). The recursive-multiple-priors model imposes the condition "rectangularity" on the set of priors, and Epstein and Schneider (2003) show that if the set of priors is rectangular, dynamic consistency is obtained. We will soon introduce the definition of rectangularity in the context of our setting, and define an ambiguity averse player and ambiguity averse Pareto efficiency.

In order to minimize the impact of the difference between the truth and a player's prior on the long run survival, we assume that every player has the largest rectangular set of priors. The reason is the following. In the recursive-multiple-priors model, an ambiguity averse player can be thought of being Bayesian. Therefore, the "distance" between the truth and a player's set of priors can be a determinant of the player's long run survival. This is indeed true in the presence of the Bayesian expected utility maximizer with the correct beliefs. Condie's result (Theorem 1, 2008) states that if there exists at least one Bayesian expected utility maximizer who knows the truth, an ambiguity averse player with a rectangular set of priors vanishes unless the ambiguity averse player has the prior sufficiently close to the truth. The argument is similar to Blume and Easley (2006). One Bayesian player knows the truth, and an ambiguity averse player with a rectangular set of priors can be thought of as a Bayesian learner. Consequently, if her choice of the worst prior is far away from the truth, the ambiguity averse player is more likely to vanish.

To isolate the impact of ambiguity aversion on long run survival, we will make the assumption that every player has the largest rectangular set of priors. And, we show that for every Pareto efficient allocation with commitment, there exists a stationary Pareto efficient allocation. This means that if every player is ambiguity averse, Pareto efficiency with commitment implies that every player's consumption could depend on the state only. This, in turn, implies that no player's consumption goes to zero.

We define an ambiguity averse player.

Definition 2.7.1. *Player i is said to be ambiguity averse*

1. *if she has a set of priors $\mathscr{P}_i$, and*

2. *if for every consumption plan c, her expected utility in period 0 is given by $\inf_{\mu_i \in \mathscr{P}_i} V_i(c, \mu_i)$.*

This definition is ex-ante ambiguity averse. That is, Definition 2.7.1 means that an ambiguity averse player is ambiguity averse before the risk sharing game begins, and becomes Bayesian after the risk sharing game begins. This raises an issue of dynamic consistency such as Ellsberg paradox.

For instance, $\mu_i \in \mathscr{P}_i$ is player i's worst prior at the outset, however, after observing the state s_1, $\mu_i(\cdot|s_1)$ may be no longer player i's worst prior.

In order to handle dynamic consistency, we adopt recursive multiple priors model in Epstein and Schneider (2003).[4] The important notion in recursive multiple priors model is *rectangularity*. We introduce the definition of rectangularity below. It is originally introduced in Epstein and Schneider (2003), and modified in the context of our setting.

Definition 2.7.2 (Epstein and Schneider)**.** *A set of priors $\mathscr{P}$ is rectangular if for every $\mu^{(0)}, \mu^{(1)}, \cdots, \mu^{(s)} \in \mathscr{P}$ and every $s^t \in \mathcal{F}_\infty$, there exists $\mu \in \mathscr{P}$ such that for all $E \subset \mathcal{S}$,*

$$\mu(E|s^t) = \sum_{s \in S} \mu^{(s)}\left(E|(s^t, s)\right) \mu^{(0)}\left((s^t, s)|s^t\right).$$

To understand Definition 2.7.2, let us imagine the following situation. The consumption plan c is given to a player who has a set of priors $\mathscr{P}$. At the state history s^t, she believes $\mu^{(0)}$ is the worst prior among $\mathscr{P}$. In the subsequent state history (s^t, s), the player chooses $\mu^{(s)}$ as the worst prior. If the player were Bayesian, her prior conditional on s^t is $\mu^{(0)}$ and her prior conditional on (s^t, s) is $\mu^{(s)}$. Rectangularity means that there exists a prior $\mu \in \mathscr{P}$ such that conditional on the state history s^t, $\mu(\cdot|s^t)$ equals $\mu^{(0)}(\cdot)$ and conditional on the state history (s^t, s), $\mu(\cdot|s^t, s)$ equals $\mu^{(s)}(\cdot)$. Therefore, if a player has a rectangular set of priors, the player can be thought of as Bayesian, updating the set of priors by applying Bayes' rule prior-by-prior.

From the previous paragraph, we can see that rectangularity implies dynamic consistency. This is because given the consumption plan c, if μ is the worst prior in period 0, $\mu(\cdot|s^t)$ is the worst prior conditional on the state

[4] For other ways of considering dynamic consistency, see Ozdenoren and Peck (2008) and Siniscalchi (2011)

history s^t.[5] This means that if player i is ambiguity averse, and if her set of priors $\mathscr{P}_i$ is rectangular, then for every consumption plan c and the state history s^t, her expected utility conditional on s^t is given by $\inf_{\mu_i \in \mathscr{P}_i} V_i(c, \mu_i | s^t)$.

It would not be uncommon that a rectangular set of priors is large. For instance, if a set of priors consists of i.i.d.processes only, it is not rectangular. We would have to allow for general stochastic processes to have a rectangular set of priors. This means that in order to achieve dynamic consistency, we would have to depart from i.i.d. processes. This departure would imply that the difference between the Bayesian case and the ambiguity aversion case is partially due to the difference between i.i.d. processes and the general true data generating processes.

We turn to Pareto efficiency when all players are ambiguity averse.

Definition 2.7.3. *A strategy profile σ is ambiguity averse Pareto efficient if there exists no strategy profile σ' such that for each $i \in I$,*

$$\inf_{\mu \in \mathscr{P}_i} V_i(c_i^{\sigma}, \mu) \leq \inf_{\mu \in \mathscr{P}_i} V_i(c_i^{\sigma'}, \mu_i)$$

and for some $i \in I$, the above inequality is a strict inequality.

Our analysis will be focused on the case where every ambiguity averse player has the largest rectangular set $\mathscr{P}^{\dagger}$ of priors. $\mathscr{P}^{\dagger}$ is the set of probability measures μ defined on $(\mathcal{S}, \mathcal{F})$ such that for all $s^t \in \mathcal{F}_{\infty}$, $\mu(s^t) > 0$.

The reason why we restrict our attention to $\mathscr{P}^{\dagger}$ is that we want to isolate the effect of ambiguity aversion on the long run survival. As we have shown in Bayesian cases, closeness between the true data generating process and players' priors plays a critical role in determining whether a player survives

[5]For precise statements, see Epstein and Schneider Theorem 3.2 and Theorem B.1 (2003)

in the long run. So, if we assume a smaller set of priors, whether a player survives in the long run is affected by not only ambiguity aversion but also how close her set of prior is to the truth. By assuming the largest rectangular set of priors, we would be able to eliminate the impact of closeness between the truth and the player's prior and single out the impact of ambiguity aversion on the long run survival. It would be a future project to study what happens when different players have different sets of priors.

Under the assumption that $\mathscr{P}_i = \mathscr{P}^\dagger$ for all $i \in I$, a stationary strategy profile, which we define below, turns out to be useful in the characterization of ambiguity averse Pareto efficiency.

Definition 2.7.4. *A strategy profile σ is stationary if for all $i \in I$, c_i^σ depends only on the state.*

When a strategy profile σ is stationary, $c_i^\sigma(s)$ refers to player i's consumption in state $s \in S$ by abusing notation. The following lemma shows that given an ambiguity averse Pareto efficient allocation, one can find a payoff-equivalent, stationary, ambiguity averse Pareto efficient allocation.

Lemma 2.7.1. *Suppose that $\mathscr{P}_i = \mathscr{P}^\dagger$ for all $i \in I$ and that a strategy profile σ is ambiguity averse Pareto efficient. Then, there exists a stationary ambiguity averse Pareto efficient strategy profile σ^* such that for all $i \in I$,*

$$\inf_{\mu \in \mathscr{P}} V_i(c_i^\sigma, \mu) = \frac{1}{1-\delta} u_i(c_i^{\sigma^*}(1))$$

.

(proof) See the Appendix.

Lemma 2.7.1 means that given an ambiguity averse Pareto efficient strategy, one can find a payoff-equivalent, stationary, ambiguity averse Pareto efficient strategy. There may exist other types of ambiguity averse Pareto efficient strategies. However, Lemma 2.7.1 would allow us to focus on stationary and ambiguity averse Pareto efficient strategies.

We sketch the proof of Lemma 2.7.1. Given the ambiguity averse Pareto efficient strategy profile σ, we can find a certainty equivalent. Let $\overline{c_i^\sigma}$ be a certainty equivalent for player i. For any ν_q, player i's expected utility in period 0 should be larger than or equal to the utility of receiving $\overline{c_i^\sigma}$ in every state and every period. The concavity of utility functions implies that $\overline{c_i^\sigma}$ is less than or equal to the present value of expected lifetime consumption. This implies that the sum of the players' certainty equivalents is less than or equal to the present value of expected lifetime aggregate endowment. The lowest value of expected lifetime aggregate endowment is obtained when Nature always chooses state 1 in every period because the aggregate endowment is the lowest in state 1. So, we can conclude that $\sum_{i \in I} \overline{c_i^\sigma} \leq e(s_1)$. This inequality should hold with equality, otherwise σ is not Pareto efficient.

By definition of stationary strategy profiles, the following proposition is immediate.

Proposition 2.7.1. *Suppose σ is a stationary, ambiguity averse Pareto efficient strategy profile. Then, for each $i \in I$, there exists $\underline{c_i}$ such that $c_i^\sigma(s^t) \geq \underline{c_i}$ for all state history s^t.*

2.8 Ambiguity Averse Players and Self Enforcement

We now discuss Pareto efficient subgame perfect equilibria when all players are ambiguity averse. Because the concept of subgame perfect equilibrium is built on the notion of time consistency, it is only well-defined in our setting if players' priors ensure that players are time consistent. Therefore, throughout this section we make the assumption that each player has a rectangular set of priors. As we discussed earlier, this implies time consistency. It is then immediate how to define subgame-perfect equilibrium. We make the assumption of rectangularity throughout this section without explicit mentioning.

We would like to mention that the players are ambiguity averse only about the true data generating process. We do not assume that the players are ambiguity averse about other players' strategies.

In this section, we provide two results. For the first result, we assume that every player has the largest rectangular set of priors. Under this assumption, the first result of this section is that a stationary, ambiguity averse Pareto efficient strategy profile with commitment can be a Pareto efficient subgame perfect equilibrium for large discount factor, provided that every player's continuation payoff is strictly larger than her expected utility of consuming her own endowment. Since we assume the largest rectangular set of priors, if the strategy profile is stationary, every player believes that she consumes the lowest amount in every period. This means that equilibrium consumption is also stationary, and therefore, for a sufficiently large discount factor, a stationary, ambiguity averse Pareto efficient strategy profile with commitment can be a subgame perfect equilibrium.

For the second result, we assume that different players have different rectangular sets of priors, but share at least one common prior, and the aggregate endowment is constant. The second result in this section is that if players are sufficiently patient, in a Pareto efficient subgame perfect equilibrium, every player consumes a constant amount of good in every period and every state.

To derive the two results, we first characterize subgame perfect equilibria. Since rectangularity guarantees dynamic consistency, the characterization of subgame perfect equilibria when players are ambiguity averse is the same as the characterization of subgame perfect equilibria when the players are Bayesian. The no-transfer equilibrium is the worst subgame perfect equilibrium, and any subgame perfect equilibrium can be sustained by a grim trigger strategy profile. After the characterization of subgame perfect equilibria, we state the two results.

We turn to characterize subgame perfect equilibria. Since rectangularity implies dynamic consistency, Lemma 2.5.1 is carried over. The no-transfer strategy profile is a subgame perfect equilibrium, and it is the worst subgame perfect equilibrium. Lemma 2.5.2 is also carried over. The unique Nash equilibrium of the static risk sharing game is that every player makes zero transfer to everyone else, and every player receives their minimax payoffs in the unique Nash equilibrium of the stage game. Therefore, we state the following Lemmas without proofs.

Lemma 2.8.1. *The "no-transfer" strategy profile* $\underline{\sigma}$ *is a subgame perfect equilibrium. And, in any subgame perfect equilibrium* σ*,* $\inf_{\mu \in \mathscr{P}_i} V_i(c_i^\sigma, \mu | s^t) \geq \inf_{\mu \in \mathscr{P}_i} V_i(e_i, \mu | s^t)$ *for every* $i \in I$ *and every* $s^t \in \mathcal{F}_\infty$*.*

Lemma 2.8.2. *For an allocation* $(c_1, \cdots, c_I)$*, there exists a subgame perfect equilibrium* σ *with* $c_i^\sigma = c_i$ *for all* i *if and only if for all* $i \in I$ *and for* $s^t \in \mathcal{F}_\infty$*,*

$$u_i(c_i(s^t)) + \delta \inf_{\mu \in \mathscr{P}_i} V_i(c_i, \mu | s^t) \geq u_i(e_i(s^t)) + \delta \inf_{\mu \in \mathscr{P}_i} V_i(e_i, \mu | s^t).$$

We turn to characterize a class of Pareto efficient subgame perfect equilibrium when the players are ambiguity averse.

Proposition 2.8.1. *Suppose that for all* $i \in I$*,* $\mathscr{P}_i = \mathscr{P}^\dagger$*. Let* σ *be a stationary, ambiguity averse Pareto efficient strategy profile such that for all* i*,* $c_i^\sigma(1) > \min_{s \in S} e_i(s)$*. There exists* $\underline{\delta}$ *such that for all* $\delta > \underline{\delta}$*,* σ *is a Pareto efficient subgame perfect equilibrium.*

Proof. In the Appendix. □

The idea for Proposition 2.8.1 is straightforward. Lemma 2.7.1 allows us to focus on stationary, ambiguity averse Pareto efficient strategy profiles, and as a result, we may not have to worry about the existence of a noted strategy profile in Proposition 2.8.1. Once we have a stationary strategy profile, clearly, for a sufficiently large discount factor, the stationary strategy profile can be a subgame perfect equilibrium. Pareto efficiency automatically follows when the stationary strategy profile is Pareto efficient with commitment.

For the following proposition, we consider the conventional environment, which can be found in Billot *et al.* (2000), Rigotti, Shannon and Strzalecki (2008), and Ghirardato and Siniscalchi (2016). The typical assumptions are 1) that different players have different sets of priors but share at least one prior, and 2) that aggregate endowment is constant across states. A strategy profile σ is called a full insurance if there exists a constant consumption plan $c = (\bar{c}_1, \cdots, \bar{c}_I)$ such that for all $i \in I$ and all $s^t \in \mathcal{F}_\infty$, $c_i^\sigma(s^t) = \bar{c}_i$.

Proposition 2.8.2. *Assume the followings.*

1. $\nu_q \in \cap \mathscr{P}_i$ *for some* q *in the interior of* ΔS.

2. $e(1) = e(2) = \cdots = e(K) \equiv e$.

Then, there exist a full insurance σ *and* $\underline{\delta} \in (0,1)$ *such that for all* $\delta \geq \underline{\delta}$, σ *is a Pareto efficient subgame perfect equilibrium.*

Note that we do not assume that $\mathscr{P}_i = \mathscr{P}^\dagger$ in Proposition 2.8.2. The idea for Proposition 2.8.2 is basically the same as Billot *et al.* (2000). Under the assumption that every player shares at least one prior and there is no uncertainty of aggregate endowment, Pareto efficiency implies a full insurance in complete markets. There are two minor differences between Billot *et al.* Theorem 1 and Proposition 2.8.2. The first difference is rectangularity. Since we consider dynamic environments, we added the condition that a set of priors is rectangular so as to update the set of priors over time. Rectangularity is unnecessary in Billot *et al.* (2000) as they consider complete markets.

The second difference is the assumption on the shared prior. The second condition in Proposition 2.8.2 means that all players share a point prior that represents an independently and identically distributed process. For Billot *et al.*'s result (Theorem 1, 2000), no assumption needs to be made on the common prior shared by all players. This is partially because Billot *et al.* (2000) considers complete markets, which means that external enforcement is available. In our model, we need to make sure that every player's continuation payoff does not converge to their full insurance payoff. If a player's minimax payoff converges to her full insurance payoff along a certain path, it would be impossible to find the lower bound for the discount factor. This means that the player's self enforcement condition would be violated at some history. By requiring that every player share a common i.i.d. process, we can guarantee that for some full insurance, every player's minimax payoff at every history does not converge to their full insurance payoffs. This, in turn, guarantees the existence of the lower bound for the discount δ.

In Proposition 2.8.1, we assume the largest rectangular set of priors, but allow for uncertainty of aggregate endowment. In Proposition 2.8.2, we allow for different rectangular sets of priors, but assume no uncertainty of aggregate endowments. It would be interesting to analyze the case in which different players have different sets of priors and there is uncertainty about aggregate endowments.

2.9 Conclusion

In this section, we summarize our results.

We analyze impacts of self enforcement and ambiguity aversion on the long run survival in an infinite horizon risk sharing game. Our first analysis characterizes a Pareto efficient allocation when all players are standard Bayesian and external enforcement is available. It is shown that player j's consumption converges to zero if 1) player i learns a true data generating process faster than player j, 2) a true data generating process belongs to the support of player i's prior, but does not belong to the support of player j's prior, or 3) player i's asymptotic posterior belief is closer to a true data generating process in terms of Kullback-Leibler divergence. The first condition implies that learning speed matters, especially when more than one player learn the truth. The second condition is intuitive; if player i learns the truth while player j does not, player i will be a winner. The third condition means that even if a player does not learn the truth, she can be a winner if her asymptotic posterior belief is closer to the truth than everyone else's asymptotic posterior belief is.

Another characterization of Pareto efficient allocation under external enforcement is that when the true data generating process is closer to someone's prior than someone else's prior, any Pareto efficient allocation under external enforcement is not a subgame perfect equilibrium of the risk sharing game. The intuition is because it is difficult to change the players' beliefs after observing sufficiently many states. Pareto efficiency determines consumptions based on ratios of unconditional probabilities. If player i's prior is closer to the true data generating process than player j's prior is, the ratio of player i's prior to player j's prior grows as time goes on. If the players have observed sufficiently large number of states, the ratio would stay large in the future regardless of state realizations and player j's consumption is small for long periods because of Pareto efficiency. So, player j deviates in such a bad case if external enforcement is not available.

Our second result is the existence of a lower bound for a player's consumption in a Pareto efficient subgame perfect equilibrium. An interesting feature of the second result is that a player's consumption can be arbitrarily close to zero in no Pareto efficient subgame perfect equilibrium. The idea behind this result is similar to the reason that a Pareto efficient allocation under external enforcement is not a subgame perfect equilibrium. If a player's consumption becomes very small, it takes a long time for her consumption to become large. Therefore, sufficiently small consumption implies a profitable deviation.

Our third analysis characterizes a Pareto efficient allocation with commitment when all players are Gilboa-Schmeidler type. We first show that for any ambiguity averse Pareto efficient allocation, there exists a payoff-equivalent stationary ambiguity averse Pareto efficient allocation. Due to this, when the players are ambiguity averse, Pareto efficiency implies that each player's consumption could depend only on the state. In this sense, ambiguity averse players insure against model uncertainty and secure a certain amount of endowments.

Our last analysis is the characterization of Pareto efficient subgame perfect equilibrium when players are ambiguity averse. Due to our third analysis, we can focus on stationary Pareto efficient allocations with commitment. Since every player's endowment is stationary, it would not surprising that if the players are sufficiently patient, a stationary Pareto efficient allocation with commitment can be a Pareto efficient subgame perfect equilibrium.

We would like to mention two possible future projects. The first future project is to study the dynamic paths of consumptions in a Pareto efficient subgame perfect equilibrium when Bayesian players have heterogeneous priors. Pareto efficient allocations with commitment are primarily determined by the players' priors and independent of endowments. However, in a subgame perfect equilibrium, endowments are important because of self enforcement conditions. If a player has a large endowment today, either her today consumption or her continuation payoff must be large in order to prevent her from deviation. This argument implies that there would be a correlation between current endowments and current consumption or between current endowments and future consumptions. If we can theoretically characterize such a correlation, we can empirically study whether actual data fits into Pareto efficient allocations with commitment or Pareto efficient subgame perfect equilibria. The second future project is to consider a risk sharing game in which whether a player is Bayesian or ambiguity averse is private information. It is not clear whether a player can learn another player's type in the long run because for instance, an ambiguity averse player can pretend to be Bayesian. It would be interesting to figure out how a Pareto efficient subgame perfect equilibrium evolves as time passes by.

Chapter 3

Screening Agents with Different Qualities of Information

3.1 Introduction

To make an informed decision, people often try to learn information from others. In many cases, the decision maker, who will here be called the "principal," may not be able to collect information herself for various reasons. For instance, she may not have enough time to get information, or she may not have a sufficient skill set to process information. Therefore, she may have an incentive to hire an agent who can collect information and pass it on to her. From the principal's perspective, the problem then becomes now to design a contract between her and the agent so as to acquire true information without paying too much.

This question can be found in many areas of life. A company which wants to extend its business by developing a new product may have to figure out its customers' preferences. The company might hire an analyst who learns customers' preferences. After the analyst finishes her research, the company would base its decision on the information provided by the analyst. Another example would be a candidate for an electoral office. A candidate usually has a group of people who do research on voters' preferences. A strategy to attract voters would be based on what the candidate learns from this group of people.

Important problems that might have to be resolved in a principal agent relationship in which the principal's objective is to gain information from the agent are: (i) the agent's preferences over the action that the principal will take based on the information might not align with the principal's preferences, so that the agent has an incentive to distort the information; (ii) the agent needs to exert effort to obtain the information and might shirk on this effort. In this paper, we shall abstract from both of these problems: the information is freely available to the agent. It is exogenous. Also, the agent has no interest per se in the principal's action.

Once we have abstracted from the two problems described in the previous paragraph, only one remains: the principal might compete with other principals for the agent's time. In other words: the agent might have an outside option, and might require the principal to pay to the agent at least as much as the outside option. We study a situation in which the outside option depends on the quality of information that the agent can provide: well-informed agents have better outside options than poorly informed agents. Whereas the agent knows whether she is well-informed or not, the principal does not.

The main objective of this paper is to study the impact of the agent's type dependent outside option on the optimal contract. A stream of literatures is focused on how to incentivize the agent to truthfully report. McCarthy (1956) and Savage (1971) study proper scoring rules. Since proper scoring rules induce the agent to truthfully report only when the agent is risk neutral, Karni (2009) and Qu (2012) describe a direct revelation mechanism for eliciting the agent's subjective probabilities when the agent is risk averse. Hossain and Okui (2013) introduce the binarized scoring rule that works irrespective of the agent's risk preference. Those literatures do not consider the agent's participation constraint. In other words, in those papers, it is always assumed that the agent participates. However, it is not uncommon that the agent has a type dependent outside option. We study a characterization of the optimal contract when the agent has a type dependent option and as a result, the agent's participation constraint may not be satisfied in some cases.

The issue of optimal contracts arises when the agent has an outside option. If the agent has zero outside option and always accepts the contract proposed by the principal, optimality would be trivial. Since the agent is only interested in monetary payments, the principal can always find an incentive compatible payment scheme to elicit the agent's posterior belief. Multiplying the payment scheme by a small number, the principal can learn the agent's true posterior belief with little payment. The papers in the previous paragraph do not assume the agent's outside option, and therefore they do not need to consider optimality. However, in this paper, the agent has a type dependent outside option, and thus, the optimal way of eliciting the agent's private information needs to be taken into account.

In the presence of the agent's type dependent outside option, the principal may not want to elicit information from every type. It is possible that the principal can be better off by rejecting some types. Suppose that the principal lowers the payment to every type by some constant amount so that some types reject the contract. The principal may not make an informed decision if the agent rejects the contract, but the principal can make an informed decision if the agent accepts the contract. From an ex-ante viewpoint, the decrease in the expected payment could be larger than the decrease in the expected benefit of the principal. Another point we would like to make is that Pareto efficiency could imply that the principal rejects some types. This can happen when the increase in the expected utility of the principal is smaller than the agent's outside option. This paper attempts to analyze the impact of the principal's utility on the optimal contract.

The principal in our model will propose a contract to the agent which specifies monetary transfers conditional on the true state and the agent's report. The agent accepts the contract if his expected utility from the contract is at least equal to his outside option. The revelation principle implies that it is without loss of generality to assume that the contract requires the agent to report to the principal both the quality of his signal, and the signal realization. We assume that the true state of the world is observable and contractable ex post, that is, the agent's payment is based on his report as well as the true state. The agent is risk neutral.

In general, of course, the "quality of information" is an ambiguous phrase. In this paper we call the quality of information the "agent's type," and we impose one assumption on the agent's type so that the quality of information is represented by a mean preserving spread of posterior distributions. That is, the posterior distribution of a higher type is a mean preserving spread of the posterior distribution of a lower type. This implies Blackwell dominance. If the agent' type is higher, and if the agent reports his information truthfully, then, the principal's expected utility is higher if the quality of information is higher.

Our first result is a characterization of incentive compatibility contracts. First, we show equivalence between incentive compatible contracts and proper scoring rules. Proper scoring rules are mechanisms that induce the agent to report the true posterior belief. In our model, the principal's main interest is to figure out the true posterior belief. Intuitively, she can ask the agent to report the posterior belief instead of his type and the signal realization. Second, we show that for any incentive compatible contract, the expected payment to the agent is an increasing and convex function of the agent's type. Concavity of the outside option means that there exist at most two threshold agent types such that the principal hires the agent only if his type is worse than the lower threshold type or better than the higher threshold type.

Our second result is a necessary condition for an optimal contract. The necessary condition is based on the convexity of the expected payment to the agent and the concavity of the agent's outside option. There are five cases concerning which agent type is hired. It can be optimal for the principal to hire 1) no agent types, 2) all agent types, 3) only high types, 4) only low types, or 5) only low and high types but no middle types. The fifth case if of special interest, in which it is optimal for the principal to admit low and high types and reject intermediate types. Obviously, the participation constraint for the threshold types is binding. Interestingly, it is shown that the participation constraint for the uninformed type is not binding. That is, in the fifth case, it is optimal for the principal to pay the uninformed agent more than his outside option. The reason is the following. In our model, the optimal incentive compatible contract is such that the expected payment can be thought of as a lump sum payment plus a bonus payment that is proportional to the quality of information. With a large lump sum payment and a small bonus payment, the principal can lower the expected payments to higher types, although she would have to admit the uninformed type. In the fifth case, reduction in the expected payment to high types outweighs increase in the expected payment to low types, meaning that the principal optimally hires low and high types only and rejects middle types.

For our third result, we impose one more assumption to prove the existence of a "linear contract", the expected payment of which is linear in the agent's type. There are two important properties of linear contracts. First, in our model, if there exists at least one linear incentive compatible contract, an optimal contract is a linear incentive compatible contract. Suppose a linear incentive compatible contract exists. Given an incentive compatible contract, if it is not a linear contract, then the principal can affine-transform the linear contract so that she admits the same agent types as the incentive compatible contract while paying less to every agent type who accepts the affine-transformed contract.

The second property of linear contracts is that the marginal value of information can be positive at the zero type. This is incompatible with Radner and Stiglitz Theorem (1984, p. 36). Obviously, it must be the case that a linear contract violates some assumptions in Radner and Stiglitz. Because of equivalence between incentive compatible contracts and proper scoring rules, a proper scoring rule that corresponds to a linear incentive compatible contract also violates some assumptions in Radner and Stiglitz. It turns out that if a proper scoring rule corresponds to a linear incentive compatible contract, the only violation is that the proper scoring rule is discontinuous at the prior. If a proper scoring rule is continuous at the prior, we can use arguments analogous to the envelope theorem. It implies that for a small change in informativeness, the agent's optimal report does not change. In other words, when the increase in informativeness is small, the agent behaves as if his optimal report is the prior. Clearly, if the agent continues to report the prior, there is no change in the expected payments to the agent. In fact, a proper scoring rule does not have to be continuous at the prior. The theorem due to McCarthy, which we will state in Section 3, means that a proper scoring rule is associated with a convex function of beliefs. The proper scoring rule being continuous at the prior implies that the corresponding convex function is differentiable at the prior. However, a convex function does not have to be differentiable at the prior.

The rest of the paper is organized as follows. In Section 2, we describe our model and assumptions. In Section 3, we discuss proper scoring rules and state a theorem that completely characterize proper scoring rules. In Section 4, we study the relationship between proper scoring rules and incentive compatible contracts and show the equivalence between them. In Section 5, we illustrate additional properties of all incentive compatible contracts. In

Section 6, we characterize a necessary condition for an optimal contract. In Section 7, we consider the special case in which the expected payment to the agent is linear in the agent's type. In Section 8, we analyze an example using our results. In Section 9 we discuss the assumptions we make in this paper. Section 10 is a conclusion. All proofs that are omitted from the main text are in the Appendix.

3.1.1 Related Literature

Jullien (2000) is related to our model in terms of participation constraint. His model is an extension of standard uninformed principal-informed agent models such as Baron and Myerson (1982), Guesnerie and Laffont (1984) and Maskin and Riley (1984). In these literatures, the agent has a constant outside option, whereas in Jullien (2000) the agent has a type-dependent outside option. Assuming that virtual surplus is quasi-concave in the agent type, Jullien (2007) shows that interior types may be excluded from trade. This is similar to our setting in the sense that in our model, the agent's outside option is concave in the agent type and in some cases, middle types reject the contract proposed by the principal.

For elicitation of forecasts, there is a large volume of literatures. McCarthy (1956) and Savage (1971) characterize proper scoring rules. Gneiting and Raftery (2007) provides a review on proper scoring rules. If the agent is risk averse, his report can be biased if a proper scoring rule is implemented. To take care of this issue, Karni (2009), Qu (2012), and Hossain and Okui (2013) study a method of eliciting the agent's posterior belief when the agent is not risk neutral. In our model, the agent is risk neutral. In addition, the principal is interested in the true posterior belief itself. As a result, it is sufficient to use proper scoring rules. However, for tractability and a technical issue, we state results in terms of incentive compatibility contracts.

Regarding the agent's information structure, Amir and Lazzati (2016) share a similar assumption with this paper even though their setting is very different from ours. They study endogenous information acquisition in Bayesian games whereas in our model, information acquisition is exogenous. Nonetheless, under the assumption that signal distribution is convex in the quality of information, Amir and Lazzati (2016) show that the expected utility of a player is convex in the quality of information. In this paper, we have a similar assumption, and show that if an incentive compatible contract is proposed, the expected payment to the agent is increasing and convex in the quality of information.

According to Amir and Lazzati (2016), Arrow (1974) and Radner (2000) would support our assumption that the agent's outside option is concave in the quality of information. In Arrow (1974) and Radner (2000), it is pointed out that information production often has increasing returns to scale, which means that the marginal cost of producing higher quality of information is decreasing. In our model, the agent can be thought of as an information supplier. Since we interpret the agent's type as the quality of information, the decreasing marginal cost of producing a higher quality of information implies that the agent's outside option is a concave function of the agent's type.

Our model provides a counter-example of Radner and Stiglitz (1984), in which it is shown that under mild assumptions, the marginal value of information is zero. It turns out that one of the assumptions in Radner and Stiglitz (1984) is violated in some circumstances. We analyze an example in which the marginal value of information is positive when the principal proposes an optimal contract.

3.2 Model

There is one principal and one agent. There are n possible states of the world, and the set of the states is $\Omega = \{\omega_1, \cdots, \omega_n\}$. The principal is uncertain about the state of the world and has a prior belief $\vec{\mu}_0 = \Big(\mu_0(\omega_1), \cdots, \mu_0(\omega_n)\Big) \in \Delta\Omega$.[6] The agent has the same prior belief $\vec{\mu}_0$ as the principal. The agent can run an experiment, and observe a signal realization.

Notations We introduce several notations that we shall use many times. For $\vec{\mu} = (\mu_1, \cdots, \mu_n) \in \mathbb{R}^n$, $|\vec{\mu}|$ denotes Euclidean norm. That is, $|\vec{\mu}|= \sqrt{\sum_{i=1}^n \mu_i^2}$. And, $\hat{\mu}$ is a corresponding unit vector.

$$\hat{\mu} = \frac{\vec{\mu}}{|\mu|}.$$

For $\vec{\mu}, \vec{\nu} \in \mathbb{R}^n$, $\vec{\mu} \cdot \vec{\nu}$ denotes the Euclidean inner product. That is, $\vec{\mu} \cdot \vec{\nu} = \sum_{i=1}^n \mu_i \nu_i$.

Integral notation will be simplified. When we denote an integral over some subset of $\mathbb{R}^k$, dy will mean $\prod_{j=1}^k dy_j$. For instance, choose a closed unit disk, D^k in $\mathbb{R}^k$. For an integrable function $f : D^k \longrightarrow \mathbb{R}$,

$$\int f(y_1, \cdots, y_k) dy_1 \cdots dy_k = \int f(y) dy.$$

$\vec{1}$ is the n-dimensional vector $(1, 1, \ldots, 1)$, and $\vec{0} = (0, \cdots, 0) \in \mathbb{R}^n$.

Timeline Nature chooses a state ω from Ω. The agent learns his type, which determines the distribution of his posterior beliefs. We will precisely describe the agent type soon. The principal proposes a contract, which specifies monetary transfers from the principal to the agent. The agent decides

[6]For a finite set W, we use ΔW to refer to the set of all measures on W.

whether to accept the contract or reject the contract. Once the agent accepts the contract, he observes a signal realization, and decides what to report. Then the principal takes an action. The true state is realized and monetary payments are made from the principal to the agent according to the contract. A true state is observable by the principal and the agent, and the true state is contractible.

The Principal The principal's utility depends on the action, the state and the monetary transfer from the principal to the agent. The principal's utility is quasi-linear in monetary transfers. A is the set of actions available to the principal. We assume that A is a compact and connected subset of $\mathbb{R}^m$ for some $m \in \mathbb{N}$. u is the principal's utility without monetary transfers:

$$u : \Omega \times A \longrightarrow \mathbb{R}$$

We assume that for every $\omega \in \Omega$, $u(\omega, a)$ is concave in a and differentiable almost everywhere with respect to a. Given the state ω, the action a and the monetary transfer t from the principal to the agent, the principal's utility is $u(\omega, a) - t$.

Since the principal is interested in the changes to her utility, we assume that u is normalized to zero when she is uninformed. That is, given the prior belief vector $\vec{\mu}_0 = (\mu_0(\omega_1), \cdots, \mu_0(\omega_n))$,

$$\sum_{\omega \in \Omega} u(\omega, a^*(\vec{\mu}_0))\mu_0(\omega) = \vec{u}(a^*(\vec{\mu}_0)) \cdot \vec{\mu}_0 = 0$$

where

$$\vec{u}(a) = (u(\omega_1, a), \cdots, u(\omega_n, a)), \text{ and}$$
$$a^*(\vec{\mu}) \in \arg\max_a \sum_{\omega \in \Omega} u(\omega, a)\mu(\omega) \;=\; \arg\max_a \; \vec{u}(a) \cdot \vec{\mu}.$$

The Agent The agent is risk-neutral, and draws his type from the interval $X = [0, 1]$. The agent type $x \in X$ determines the distribution of posterior beliefs, and it can be thought of as the quality of information. We will see later why x can be thought of as the quality of information. $P(x)$ is a cumulative distribution function of x, and it admits a differentiable density function $p(x)$ with respect to the Lebesgue measure. We assume that $\text{supp}(p) = X$ and that for all $x \in X$, $p(x) > 0$. Agent types and states of the world are independently distributed.

The agent can observe a signal realization $y \in Y$. Y is the k-dimensional sphere with radius r centered at the origin. We assume that r is appropriately chosen so that Y has volume of 1 with respect to the Lebesgue measure. For each $x \in X$ and $\omega \in \Omega$, $F(\cdot|x, \omega)$ is a probability measure on sigma algebra on Y, and F admits a differentiable density function $f(\cdot|x, \omega)$ with respect to the Lebesgue measure. We assume that for all $x \in X$, $f(y|x) = \sum_{\omega \in \Omega} f(y|x, \omega)\mu_0(\omega) = 1$.[7] Also, we assume that for all $\omega \in \Omega$ and $x \in X$, $f(\cdot|x, \omega)$ has the full support. Finally, the agent type x has the outside option $z(x)$.

Given the agent type x and the signal realization y, $\mu(\omega|x, y)$ denotes the probability of ω conditional on (x, y). That is,

$$\mu(\omega|x, y) = \frac{f(y|x, \omega)\mu_0(\omega)p(x)}{\sum_{\omega \in \Omega} f(y|x, \omega)\mu_0(\omega)p(x)}.$$

[7] This is without loss of generality. One can always satisfy this assumption by redefining y.

The posterior belief conditional on (x, y) is oftentimes expressed as a vector.

$$\vec{\mu}(x, y) = \Big(\mu(\omega_1|x, y), \cdots, \mu(\omega_n|x, y)\Big).$$

Assumptions We impose the following assumption on f.

Assumption 1. There exist $\eta_1 : X \times Y \longrightarrow [0, \infty)$ and $\eta_2 : \Omega \times Y \longrightarrow \mathbb{R}$ such that for all ω, x, y,

$$f(y|x, \omega) = 1 + \eta_1(x, y)\eta_2(\omega, y).$$

In addition, $\eta_1(x, y)$ is increasing and convex in x for every $y \in Y$, and $\eta_1(0, y) = 0$ for all $y \in Y$.

The functions $\eta_1(x, y)$ and $\eta_2(\omega, y)$ have the following interpretation: η_2 indicates the direction into which the likelihood of signal realization y is changed when the true state is ω, and η_1 indicates the extent to which it is shifted into this direction when the agent's type is x. The more the likelihood is shifted, the more informative the signal is. Therefore, the assumption that η_1 is increasing in x means that the informativeness of the signal increases in the agent's type. The assumption that η_1 is convex means that the "marginal increase in informativeness" is the larger the larger x is. In some cases, this assumption is without loss of generality, because one can re-define x appropriately, but in general, the assumption is restrictive, and it becomes more restrictive in combination with Assumption 2 below. We postpone a discussion of these assumptions until Assumption 2 has also been introduced.

For $f(y|x, \omega)$ to be a well-defined density, η_1 and η_2 have to satisfy some conditions. Since f is a density function, it must be true that for all x and ω, $f(y|\omega, x)$ is non-negative. That is, $\eta_1(x, y)\eta_2(\omega, y) \geq -1$. And, it must be also true that

$$\int \eta_1(x,y)\eta_2(\omega,y)dy = 0.$$

This means that $\int f(y|x,\omega)dy = 1$ for all ω and x. For technical convenience, we assume that $\eta_1(x,y)$ is bounded and differentiable everywhere and that for every $\omega \in \Omega$, $\eta_2(\omega,y)$ is bounded and differentiable with respect to y.

We provide one example of a function f. $\eta_1(x,y)$ has spherical symmetry in y for every x. That is, for every $x \in X$ and $y \in Y$, $\eta_1(x,y) = \eta_1(x,-y)$. $\eta_2(\omega,y)$ is anti-symmetric in y given ω. That is, for every $\omega \in \Omega$ and $y \in Y$, $\eta_2(\omega,y) = -\eta_2(\omega,-y)$. And, there exists $m \in (0,1)$ such that $-m \leq \eta_1(x,y)\eta_2(\omega,y) \leq m$ for every ω, x and y. Since $m \in (0,1)$, $f(y|x,\omega) > 0$ for all x, all y and ω. Also, since $\eta_1(x,y)\eta_2(\omega,y)$ is anti-symmetric in y for every x and every ω, $\int \eta_1(x,y)\eta_2(\omega,y)dy = 0$ for all x and ω.

It is instructive and useful to describe the implications of Assumption 1 for posterior beliefs. Let us calculate the posterior belief conditional on (x,y). Since $\sum_{\omega\in\Omega} f(y|x,\omega)\mu_0(\omega) = 1$,

$$\mu(\omega|x,y) = f(y|x,\omega)\mu_0(\omega) = \mu_0(\omega)(1+\eta_1(x,y)\eta_2(\omega,y)).$$

Define for each ω,

$$\lambda(\omega,y) = \mu_0(\omega)\eta_2(\omega,y).$$

Using vector notation,

$$\vec{\lambda}(y) = \Big(\mu_0(\omega_1)\eta_2(\omega_1,y), \cdots, \mu_0(\omega_n)\eta_2(\omega_n,y)\Big).$$

Then, $\vec{\mu}(x,y) = \vec{\mu}_0 + \eta_1(x,y)\vec{\lambda}(y)$. Define

$$\xi(x,y) = \eta_1(x,y)\left|\vec{\lambda}(y)\right|.$$

Using ξ and $\hat{\lambda}$,

$$\vec{\mu}(x,y) - \vec{\mu}_0 = \xi(x,y)\hat{\lambda}(y).$$

Recall that $\hat{\lambda}(y)$ is the normalized vector λy. Therefore:

$$|\vec{\mu}(x,y) - \vec{\mu}_0| = \xi(x,y)|\hat{\lambda}(y)| = \xi(x,y),$$

i.e. $\xi(x,y)$ is the Euclidean distance between the prior and the posterior. Note that ξ has the same property as η_1. $\xi(x,y)$ is increasing and convex in x for every y. Also, $\xi(0,y) = 0$ for every y. Finally, ξ is differentiable. Figure 3.1 shows an example when there are three states and Y is similar to a circle.

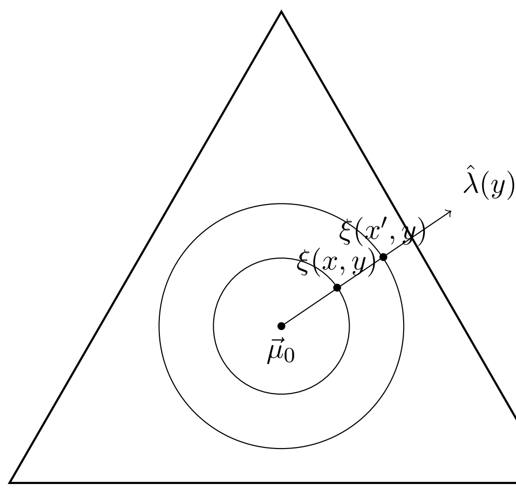

Figure 3.1: Geometric presentation of ξ and $\hat{\lambda}$

In Figure 3.1, the triangle is $\Delta\Omega$. The inner circle is the set of possible posterior beliefs that the agent of type x have. The outer circle is the set of possible posterior beliefs that the agent of type $x' > x$ have. $\hat{\lambda}(y)$ is the direction in which the posterior belief conditional on (x,y) changes as x increases. Given y, $\xi(x,y)$ the Euclidean distance between the prior belief and the posterior belief conditional on (x, y).

Figure 3.1 helps us to illustrate two points. The first point is that Assumption 1 implies the mean preserving spread of the posterior distribution. In other words, the posterior distribution of a higher type is a mean preserving spread of the posterior distribution of a lower type. The second point is that the principal would be able to back out the agent's type and the signal realization from the agent's posterior belief. That is, reporting the posterior belief is equivalent to reporting the agent's type and the signal realization.

We would like to discuss ξ and $\hat{\lambda}$ because we will use them many times instead of η_1 and η_2 in this paper. $\xi(x, y)$ determines the Euclidean distance between the prior belief and the posterior belief conditional on (x, y). Since $\hat{\lambda}(y)$ is a vector with a unit Euclidean norm, it determines a change in direction from the prior belief vector to the posterior belief vector given (x, y). The assumed monotonicity of η_1 in x implies that if the agent type is low, the corresponding posterior beliefs are closer to the prior belief. When $x = 0$, the agent is uninformed since his posterior belief is always $\vec{\mu}_0$.

We impose the following assumption on the outside option $z(x)$.

Assumption 2. The outside option $z(x)$ is differentiable, strictly increasing, and strictly concave in x.

If the agent has zero outside option, then the principal can elicit the agent's posterior belief with arbitrarily small costs. Sharing the principal's profit is one way to elicit the agent's posterior belief without having to incur large costs. The principal can sell an arbitrarily small fraction of her utility. Clearly, the agent has no incentive to report a false posterior belief, and the principal pays little amount of money to the agent. In this sense, a non-zero outside option is necessary to make the model non-trivial.

The combination of Assumptions 1 and 2 is quite restrictive. Those two assumptions basically mean that there exists a parameterization such that for every $y \in Y$, $\xi(x, y)$ is increasing convex in x and $z(x)$ is concave in x. That is, the better informed the agent is, the larger is the "informational value" of being slightly better informed in the environment described in our model, but the smaller is the increase in the outside option that results from being slightly better informed. This might be plausible in a situation in which the agent has a skill in information acquisition that is particularly useful for the principal for whom the agent works, but that is less useful outside of the relation. The analysis in this paper without Assumption 2 would require significant additional work, but some of the ideas and techniques in the paper also apply if Assumption 2 is not made.

We discuss the two assumptions in the first subsection of the Appendix because we would like to describe what happens to results if the two assumptions are violated.

The Principal's Objective The principal offers a contract to the agent. We restrict our attention to a direct mechanism. A contract h will depend on the realized state ω and the agent's report (x, y). That is,

$$h : \Omega \times X \times Y \longrightarrow \mathbb{R}.$$

In this work, we assume that for every $\omega \in \Omega$, a contract $h(\omega, x, y)$ is differentiable almost everywhere with respect to (x, y).

The principal's objective is to design an optimal contract in order to to maximize her net benefits. Given a contract h, let $\widetilde{X}(h)$ be the set of agent types who accept the contract h. Let

$$\vec{h}(x, y) = (h(\omega_1, x, y), \cdots, h(\omega_n, x, y)).$$

The principal chooses a contract h and a set $\widetilde{X}(h)$ of the agent's types who accept h to maximize:

$$\int_{\widetilde{X}(h)} p(x)dx \int_{Y} dy \quad \left[\vec{u}(a^*(\vec{\mu}(x,y))) - \vec{h}(x,y)\right] \cdot \vec{\mu}(x,y)$$

subject to the incentive compatibility condition

$$\vec{h}(x,y) \cdot \vec{\mu}(x,y) \geq \vec{h}(x',y') \cdot \vec{\mu}(x,y) \quad \forall x \in X, \forall x' \in X, \forall y \in Y, \forall y' \in Y$$

and the agent's participation constraint

$$\int \vec{h}(x,y) \cdot \vec{\mu}(x,y) dy \geq z(x) \quad \text{if and only if} \quad x \in \widetilde{X}(h).$$

A contract h is said to be incentive compatible if it satisfies the incentive compatibility condition.

The Agent's Objective After learning his type, the agent determines whether to accept or reject the contract proposed by the principal. Since the agent is risk neutral, he would accept the contract if the expected payments are higher than his outside option. Otherwise, the agent would reject the contract.

3.3 Proper Scoring Rules

There is a close connection between incentive compatible contracts and what is known in the statistics literature as "proper scoring rules." Proper scoring rules are incentive schemes that are designed to provide agents with incentives to correctly reveal their beliefs. Since in our setting for the principal it is sufficient to learn the agent's posterior beliefs, one can re-formulate the problem of designing an incentive compatible contract as the problem of designing a proper scoring rule. It turns out that some details complicate this at first sight straightforward connection between proper scoring rules and incentive compatible contracts. In this section, we describe proper scoring rules and introduce a theorem that completely characterizes proper scoring rules. In the next section we then discuss the connection between proper scoring rules and incentive compatible contracts.

We build on the definitions and terminology used in Gneiting and Raftery (2007), and modify them for the context of our setting. In this section, $\mathcal{P}$ is a convex subset of $\Delta\Omega$ and $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$.

Definition 3.3.1. *A scoring rule S is an extended real-valued function:*

$$S : \Omega \times \mathcal{P} \longrightarrow \overline{\mathbb{R}}.$$

According to Definition 3.3.1, a scoring rule S takes a state and a posterior belief, and returns a "score." Given a scoring rule S, $\vec{S}$ denotes the vector version of S. That is,

$$\vec{S}(\vec{\mu}) = \Big(S(\omega_1, \vec{\mu}), \cdots, S(\omega_n, \vec{\mu})\Big).$$

Definition 3.3.2. *A scoring rule S is regular relative to $\mathcal{P}$ if $\vec{S}(\vec{\mu}) \cdot \vec{\mu}'$ is real-valued for all $\vec{\mu}, \vec{\mu}' \in \mathcal{P}$, except possibly that $\vec{S}(\vec{\mu}) \cdot \vec{\mu}' = -\infty$ if $\vec{\mu} \neq \vec{\mu}'$.*

Definition 3.3.2 implies that given a regular scoring rule, the expected score is finite when the true posterior belief is reported. And, it is possible that for some posterior belief, the agent expects the negative infinity score if he reports a false posterior belief.

Definition 3.3.3. *A scoring rule S is proper relative to $\mathcal{P}$ if for all $\vec{\mu}, \vec{\mu}' \in \mathcal{P}$,*

$$\vec{S}(\vec{\mu}) \cdot \vec{\mu} \geq \vec{S}(\vec{\mu}') \cdot \vec{\mu}.$$

Regarding Definition 3.3.3, there is the implicit assumption that the agent prefers a higher expected score to a lower expected score. Under this assumption, Definition 3.3.3 means that when a proper scoring rule is proposed, the agent has no incentive to report a false posterior belief. This property corresponds to the incentive compatibility condition.

Now, we introduce a useful theorem. The theorem below is originally due to McCarthy (1956). We state the version provided in Gneiting and Raftery (2007, Theorem 1, p. 361). That version uses the concept of a "subgradient". For a convex function $\phi : \mathcal{P} \longrightarrow \mathbb{R}$, a function $\vec{G} : \mathcal{P} \longrightarrow \mathbb{R}^n$ is called a subgradient of ϕ if for all $\vec{\mu}_0, \vec{\mu} \in \mathcal{P}$:

$$\phi(\vec{\mu}) - \phi(\vec{\mu}_0) \geq \vec{G}(\vec{\mu}_0) \cdot (\vec{\mu} - \vec{\mu}_0).$$

If a convex function ϕ is differentiable in the interior of $\mathcal{P}$, $\vec{G} = \nabla\phi$.

Theorem (McCarthy). *A regular scoring rule S is proper relative to $\mathcal{P}$ if and only if there is a convex function $\phi : \mathcal{P} \longrightarrow \mathbb{R}$ with subgradient $\vec{G}$ such that for every $\vec{\mu} \in \mathcal{P}$:*

$$\vec{S}(\vec{\mu}) = \left(\phi(\vec{\mu}) - \vec{G}(\vec{\mu}) \cdot \vec{\mu}\right) \vec{1} + \vec{G}(\vec{\mu}).$$

The theorem is the complete characterization of proper scoring rules. Let us briefly discuss the theorem. If S is a proper scoring rule relative to $\mathcal{P}$, the expected score $\vec{S}(\vec{\mu}) \cdot \vec{\mu}$ equals $\phi(\vec{\mu})$, as one can easily calculate, and is therefore a convex function from $\mathcal{P}$ to $\mathbb{R}$. Conversely, if ϕ is a convex function from $\mathcal{P}$ to $\mathbb{R}$, then one can construct a proper scoring rule relative to $\mathcal{P}$ using ϕ and a subgradient of ϕ.

3.4 Relationship Between Proper Scoring Rules and Incentive Compatible Contracts

In this section, we study the relationship between proper scoring rules and incentive compatible contracts. Let $M = \{\vec{\mu}(x,y) \in \Delta\Omega | (x,y) \in X \times Y\}$ and M^{co} be the convex hull of M.

We raise two questions. The first question is: given a proper scoring rule, is there a corresponding incentive compatible contract? The second question is: given an incentive compatible contract, is there a corresponding proper scoring rule? The answer to the first question is clear since a proper scoring rule can be an incentive compatible contract. However, the answer to the second question is not obvious because incentive compatible contracts and proper scoring rules are defined on different domains. Fortunately, the answer to the second question is "yes".

Proposition 3.4.1. *The followings hold.*

1. *Given a proper scoring rule S relative to M^{co}, there exists an incentive compatible contract h such that $(x,y) \in X \times Y$:*

$$\vec{h}(x,y) = \vec{S}(\vec{\mu}(x,y)).$$

2. *Given an incentive compatible contract h, there exists a proper scoring rule S relative to M^{co} such that for all $(x,y) \in X \times Y$:*

$$\vec{S}(\vec{\mu}(x,y)) \cdot \vec{\mu}(x,y) = \vec{h}(x,y) \cdot \vec{\mu}(x,y).$$

Part 1 in the proposition means that given a proper scoring rule S, there exists an incentive compatible contract h that specifies the same ex-post transfers, i.e. transfers conditional on the true state, to the agent as the proper scoring rule S assigns to the agent's posterior belief. In contract, part 2 in Proposition 3.4.1 does not necessarily mean that given an incentive compatible contract h, one can always find a proper scoring rule S that specifies the same ex-port transfers to the agent. Rather, part 2 means that given an incentive compatible contract, one can find a proper scoring rule that specifies the same "interim" expected payment to the agent as the incentive compatible contract. Since the principal designs a contract at the outset, she is indifferent between an incentive compatible contract and a proper scoring rule if she incurs the same expected costs. Part 1 and Part 2 together imply that Proposition 3.4.1 is a complete characterization of incentive compatible contracts at the "interim" stage.

Proof. (Part 1) Suppose S is a proper scoring rule relative to M^{co}. For each $(x, y) \in X \times Y$, define $\vec{h}(x, y) = \vec{S}(\vec{\mu}(x, y))$. We need to show that h is incentive compatible. For all x, y, x', y',

$$\begin{aligned}\vec{h}(x,y) \cdot \vec{\mu}(x,y) &= \vec{S}(\vec{\mu}(x,y)) \cdot \vec{\mu}(x,y) \\ &\geq \vec{S}(\vec{\mu}(x',y')) \cdot \vec{\mu}(x,y) \\ &= \vec{h}(x',y') \cdot \vec{\mu}(x,y).\end{aligned}$$

(Part 2) Suppose h is an incentive compatible contract. We first note that expected payments to the agent can be a function of posterior beliefs only. The reason is the following. Suppose that $\vec{\mu}(x, y) = \vec{\mu}(x', y')$ for some x, y, x', y'. If $\vec{h}(x, y) \cdot \vec{\mu}(x, y) > \vec{h}(x', y') \cdot \vec{\mu}(x', y')$, then the agent would not report (x', y'). So, it should be the case that $\vec{h}(x, y) \cdot \vec{\mu}(x, y) = \vec{h}(x', y') \cdot \vec{\mu}(x', y')$ if $\vec{\mu}(x, y) = \vec{\mu}(x', y')$. It is possible that $\vec{h}(x, y) \neq \vec{h}(x', y')$. However, the principal can safely replace $\vec{h}(x', y')$ with $\vec{h}(x, y)$, and that replacement has no impact on the incentive compatibility condition and the expected payments to the agent.

The above paragraph implies that given an incentive compatible contract h, one can find a function $S_h : \Omega \times M \longrightarrow \mathbb{R}$ such that for all $(x, y) \in X \times Y$,

$$\vec{S}_h(\vec{\mu}(x, y)) \cdot \vec{\mu}(x, y) = \vec{h}(x, y) \cdot \vec{\mu}(x, y).$$

This also means that there exists $\phi_h : M \longrightarrow \mathbb{R}$ such that $\phi_h(\vec{\mu}(x, y)) = \vec{h}(x, y) \cdot \vec{\mu}(x, y)$ for every $(x, y) \in X \times Y$. The following claim is the first step to prove Proposition 3.4.1.

Claim. ϕ_h is bounded and convex.

Here, we need a different notion of convexity since M does not have to be a convex subset of $\Delta\Omega$. The following definition of convexity is taken from Peters and Wakker (1987).

Definition 3.4.1. A function $\phi : M \longrightarrow \mathbb{R}$ is convex if for all convex combinations $\sum_{j=1}^{J} \alpha_j \vec{\mu}_j$ of elements of $\vec{\mu}_j$ of M, whenever $\sum_{j=1}^{J} \alpha_j \vec{\mu}_j$ is in M, we have

$$\sum_{j=1}^{J} \alpha_j \phi(\vec{\mu}_j) \geq \phi\left(\sum_{j=1}^{J} \alpha_j \vec{\mu}_j\right).$$

Note that this is the standard definition of convexity, except it has been modified by inserting the "whenever ..." phrase, so that it can be applied to functions without convex domains.

Now we prove the Claim. Boundedness is clear. Suppose $\phi_h(\vec{\mu}) = \infty$ for some $\vec{\mu} \in M$. This means that for some ω, x and y, $h(\omega, x, y) = \infty$. Therefore, whenever the agent believes ω occurs with some positive probability, he would report (x, y), which is a violation of the incentive compatibility condition. If $\phi_h(\vec{\mu}) = -\infty$ for some $\vec{\mu}(x, y) \in M$, then (x, y) would not be reported.

Convexity is a direct result of the incentive compatibility condition. Suppose that $\vec{\mu}(x,y) = \sum_{j=1}^{J} \alpha_j \vec{\mu}(x_j, y_j)$ with $(x_j, y_j) \in X \times Y$.

$$\begin{aligned}
\phi_h\left(\sum_{j=1}^{J} \alpha_j \vec{\mu}(x_j, y_j)\right) &= \sum_{j=1}^{J} \alpha_j \vec{h}(x,y) \cdot \vec{\mu}(x_j, y_j) \\
&\leq \sum_{j=1}^{J} \alpha_j \vec{h}(x_j, y_j) \cdot \vec{\mu}(x_j, y_j) \\
&= \sum_{j=1}^{J} \alpha_j \phi_h(\vec{\mu}(x_j, y_j))
\end{aligned}$$

This completes the proof of the Claim, and now we are ready to prove Proposition 3.4.1. Claim shows that ϕ_h is a convex function defined on M. So, ϕ_h satisfies the assumption in Theorem 1, Peters and Wakker (1987). Therefore, Peters and Wakker (1987) Theorem 1 implies that there exists a convex function $\tilde{\phi}_h : M^{co} \longrightarrow \mathbb{R}$ that extends ϕ_h. Their construction is the following:

$$\tilde{\phi}_h(\vec{\mu}) = \inf\left\{\sum_{j=1}^{J} \alpha_j \phi_h(\vec{\mu}(x_j, y_j)) \middle| \vec{\mu} = \sum_{j=1}^{J} \alpha_j \vec{\mu}(x_j, y_j)\right\}.$$

Since ϕ_h is bounded, $\tilde{\phi}_h$ is bounded. The theorem due to McCarthy that we cited in Section 3 implies that there exists a proper scoring rule S relative to M^{co} such that $\vec{S}(\vec{\mu}) \cdot \vec{\mu} = \tilde{\phi}_h(\vec{\mu})$ for all $\vec{\mu} \in M^{co}$. Since $\tilde{\phi}_h|_M = \phi_h$, for all $(x,y) \in X \times Y$, $\vec{S}(\vec{\mu}(x,y)) \cdot \vec{\mu}(x,y) = \tilde{\phi}_h(\vec{\mu}(x,y)) = \phi_h(\vec{\mu}(x,y)) = \vec{h}(x,y) \cdot \vec{\mu}(x,y)$. $\square$

We summarize the arguments above as follows: Instead of writing a contract $\vec{h}(x,y)$ that asks the agent to report x and y, and that promises a reward conditional on x, y, and the observed true state ω, the principal could equivalently write a contract that asks the agent to report the posterior belief $\vec{\mu}(x, y)$, and that promises a reward conditional on this posterior belief and the true state ω. This is because in our model for the principal only the posterior probabilities of the true state matter, the principal is not interested in the agent's type x and the signal realization y per se.

3.5 Further Properties of Incentive Compatible Contracts

In this section, we introduce some further properties that incentive compatible contracts have to have in our setting. We would like to mention that the proof of Proposition 3.4.1 does not rely on Assumption 1 and Assumption 2, whereas the results presented in this section need Assumption 1 and Assumption 2.

We study two properties of all incentive compatible contracts.

Lemma 3.5.1. *Suppose a contract h is incentive compatible. Then the expected payment to the agent of type x,*

$$\mathbb{E}[h|x] = \int \vec{h}(x, y) \cdot \vec{\mu}(x, y) dy$$

is increasing and convex in x.

The proof of this result is in the Appendix.

One particular incentive compatible contract is that the principal "sells the project to the agent," i.e. makes payments to the agent that exactly equal her own utility from the decision problem. Therefore, Lemma 3.5.1 also implies that the principal's utility would be increasing and convex in x if she made zero payments to the agent, yet the agent reported his type truthfully. In this sense, a higher type can provide the principal with a better quality of information.

Using Lemma 3.5.1, we can derive the following lemma.

Lemma 3.5.2. *Suppose that a contract h is incentive compatible. Then, there exist at most two types $\underline{x}$ and $\overline{x}$ with $\underline{x} \leq \overline{x}$ such that all agent types between $\underline{x}$ and $\overline{x}$ reject the contract and other types accept the contract.*

Figure 3.2 illustrates the idea behind Lemma 3.5.2. In Figure 3.2 , h is an incentive compatible contract. Lemma 3.5.1 implies that the expected payment to the agent of type x, $\mathbb{E}[h|x]$ is a convex function of x. $z(x)$ is the outside option, and it is a concave function of x by Assumption 2.

As we can see from Figure 3.2 , the expected payments are higher than the outside option only if the agent type is either below $\underline{x}$ or above $\overline{x}$. Since the agent is risk neutral, he accepts the contract only if his type is lower than $\underline{x}$ or higher than $\overline{x}$.

3.6 Optimal Contracts

In this section, we provide a necessary condition for an optimal contract. The idea on which this necessary condition is based is that a positive affine transformation of an incentive compatible contract results in another incentive compatible contract, i.e. that for $b > 0$ and $c \in \mathbb{R}$, if a contract h is incentive compatible, then $bh + c$ is also an incentive compatible contract. This is obvious. The necessary condition in this section is derived from the observation that if h is optimal, then the expected utility of the principal should not change when h is subjected to an "infinitesimally small" positive affine transformation.

Given an incentive compatible contract h, a choice of b and c pins down the threshold types $\underline{x}$ and $\overline{x}$ mentioned in Lemma 3.5.2 for acceptance of the

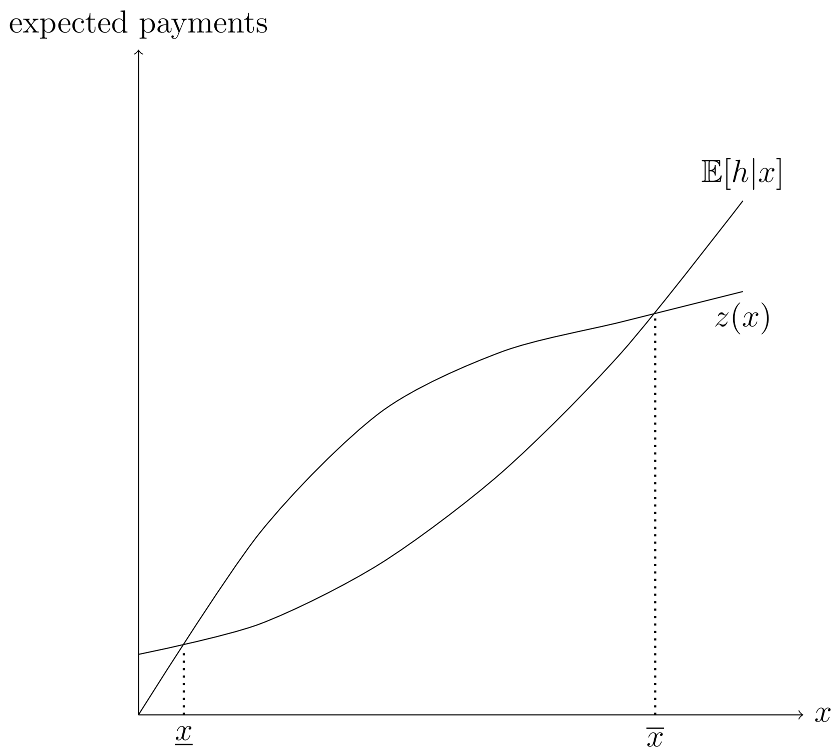

Figure 3.2: Graphical illustration of Lemma 3.5.2

contract $bh + c$. Therefore, it is possible to express the necessary condition that is based on the idea in the previous paragraph in terms of $\underline{x}$ and $\overline{x}$. This is what we do in the following lemma.

Proposition 3.6.1. *Suppose an incentive compatible contract h is optimal. Then h satisfies one of following cases.*

1. *No agent accepts h and $\mathbb{E}[h|x] < z(x)$ for all $x \in X$.*

2. *All agents accept h, and there exists a unique agent type x^* who is indifferent between accepting and rejecting. In this case $z(x^*) = \mathbb{E}[h|x^*] = \mathbb{E}[h]$.*

3. *There is a type $\overline{x} < 1$ such that only types in the interval $[\overline{x}, 1]$ accept h. In this case the threshold type and the zero type are indifferent between accepting and rejecting the contract, that is:*

$$\mathbb{E}[h|x = \overline{x}] = z(\overline{x}) \text{ and } \mathbb{E}[h|x = 0] = z(0),$$

and the threshold $\overline{x}$ satisfies the equation below:

$$\frac{z(\overline{x}) - z(0)}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{dz(\overline{x})}{dx}}(\mathbb{E}[u|\overline{x}] - z(\overline{x}))p(\overline{x}) = \mathbb{E}[h - z(0)|x \geq \overline{x}](1 - P(\overline{x})).$$

4. *There is a type $\underline{x} > 0$ such that only types in the interval $[0, \underline{x}]$ accept h. In this case the the threshold type and the highest type $x = 1$ are indifferent between accepting and rejecting the contract, that is:*

$$\mathbb{E}[h|x = \overline{x}] = z(\overline{x}) \text{ and } \mathbb{E}[h|x = 1] = z(1),$$

and the threshold $\underline{x}$ satisfies the equation below:

$$\frac{z(\underline{x}) - z(1)}{\frac{d}{dx}\mathbb{E}[h|\underline{x}] - \frac{dz(\underline{x})}{dx}}(\mathbb{E}[u|\underline{x}] - z(\underline{x}))p(\underline{x}) = \mathbb{E}[h - z(1)|x \leq \underline{x}]P(\underline{x}).$$

5. *There are types $0 < \underline{x} < \overline{x} < 1$ such that only types in the intervals $[0, \underline{x}]$ and $[\overline{x}, 1]$ accept h. In this case the threshold types indifferent between accepting and rejecting the contract, that is:*

$$\mathbb{E}[h|x = \underline{x}] = z(\underline{x}) \text{ and } \mathbb{E}[h|x = \overline{x}] = z(\overline{x}),$$

and the two thresholds satisfy the equations below:

$$\frac{z(\overline{x}) - z(\underline{x})}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{dz(\overline{x})}{dx}}(\mathbb{E}[u|\overline{x}] - z(\overline{x}))p(\overline{x}) \\ = \mathbb{E}[h - z(\underline{x})|x \leq \underline{x}]P(\underline{x}) + \mathbb{E}[h - z(\underline{x})|x \geq \overline{x}](1 - P(\overline{x})) \tag{3.1}$$

$$-\frac{z(\underline{x}) - z(\overline{x})}{\frac{d}{dx}\mathbb{E}[h|\underline{x}] - \frac{dz(\underline{x})}{dx}}(\mathbb{E}[u|\underline{x}] - z(\underline{x}))p(\underline{x}) \\ = \mathbb{E}[h - z(\overline{x})|x \leq \underline{x}]P(\underline{x}) + \mathbb{E}[h - z(\overline{x})|x \geq \overline{x}](1 - P(\overline{x})) \tag{3.2}$$

The first case is the case in which the principal's outside option is optimal. In the second case a contract that is accepted by all types is optimal. But a contract in which $\mathbb{E}[h|x]$ is always *strictly* larger than $z(x)$ is suboptimal. The principal can subtract a constant amount from the payment until the expected payment becomes tangent to the outside option curve. Thus, if a contract is optimal that all types accept, then there must be a type that is indifferent between accepting and rejecting. Note that, unlike in standard incentive problems, this need not be the lowest or highest type.

In the third case, it is optimal for the principal to hire the agent if his type is at least as good as $\overline{x}$. For this case, Proposition 3.6.1 provides two necessary conditions for an optimal contract. The first condition indicates which participation constraints have to be binding. It is not surprising that the threshold type's participation constraint has to be binding. It is somewhat more surprising that also the zero type's participation constraint has to be binding. To see why the uninformed type is indifferent, let us imagine that $\mathbb{E}[h|x = 0] < z(0)$. The principal can lower the expected payment by subjecting $\mathbb{E}[h|x]$ to an affine transformation that leaves the threshold type $\overline{x}$ unchanged, and she can do this until $\mathbb{E}[h|x = 0] = z(0)$. Figure 3.3 illustrates this affine transformation. The expected payment conditional on the type in the original contract is shown as an unbroken line, and expected payment conditional on the type in the transformed contract is shown as a dashed line. As the above graph illustrates, the principal can pay less without changing the interval of participating types.

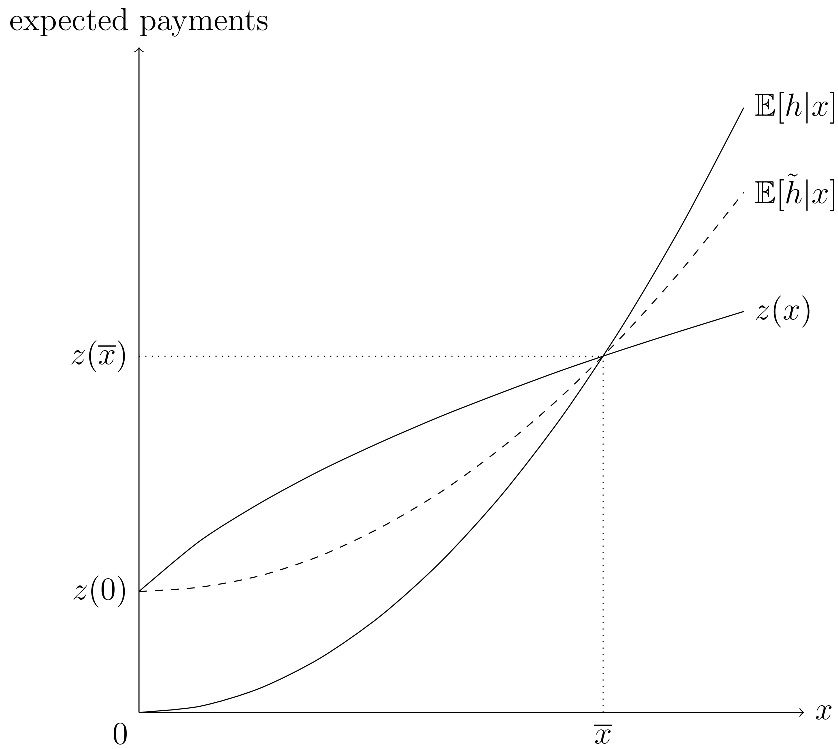

Figure 3.3: Affine transformation that makes the participation constraint of the uninformed type binding.

The second condition in the third case is a first order condition for the threshold type. Consider an incentive compatible contract h with $\mathbb{E}[h|x = 0] = z(0)$, and choose a small $\epsilon > 0$. The principal can increase the payment by ϵ percent and subtract $\epsilon z(0)$ from the payment. That is, a new contract $h' = (1 + \epsilon)h - \epsilon z(0) = h + \epsilon(h - z(0))$. This affine transformation does not affect the zero type's participation constraint, that is: $\mathbb{E}[h'|x = 0] = 0$. However, that affine transformation will lead to a change in $\overline{x}$. The necessary condition in the third case reflects that an "ϵ-affine transformation" that leads to an "infinitesimally small" change in $\overline{x}$ must not increase the expected benefit of the principal. Figure 3.4 shows the affine transformation we described in this paragraph.

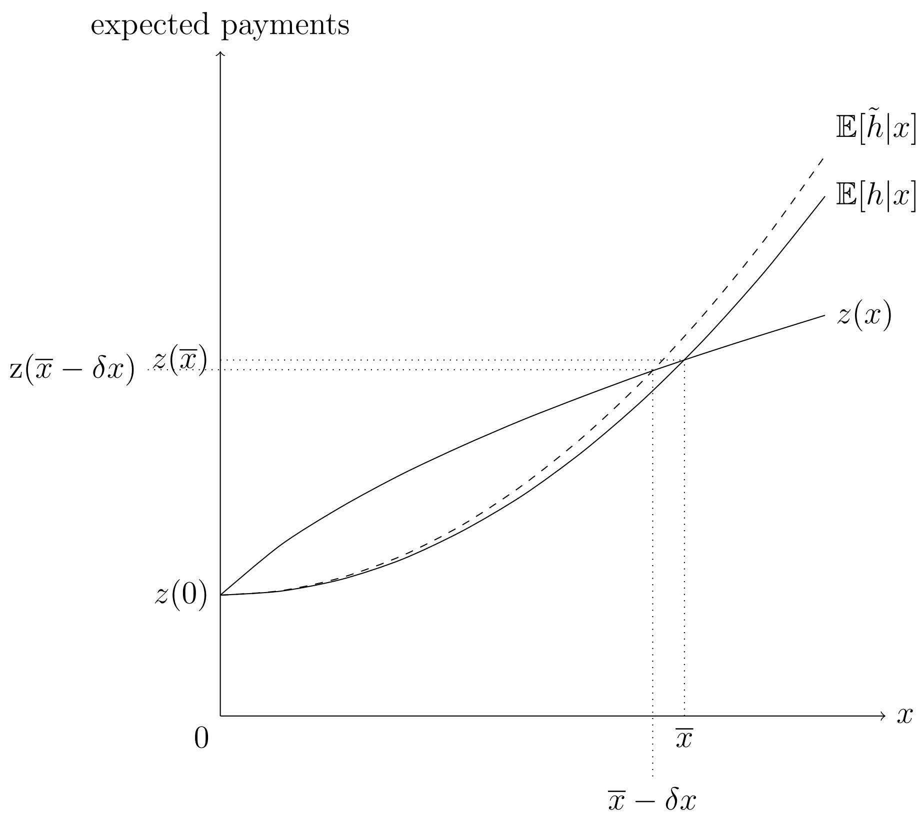

Figure 3.4: A small affine transformation while the uninformed type remains indifferent between accept and rejecting

When the threshold $\overline{x}$ changes due to the affine transformation, there are two effects. First, the principal pays more to agent types who already accept h, and this corresponds to the marginal cost. Second, new types accept the affine-transformed contract. This enables the principal to make better decisions, and therefore this corresponds to the marginal benefit. The necessary condition in Proposition 3.6.1 says that marginal cost has to equal marginal benefit.

Let us explain the marginal cost first. Up to the first order, the marginal cost is the increase in the payment times the probability of the agent accepting the contract h. Since the increase in the payment is $\epsilon(h - z(0))$ and all types higher than or equal to $\overline{x}$ accept h, the marginal cost is $\epsilon\mathbb{E}[h - z(0)|x \geq \overline{x}](1 - P(\overline{x}))$. This is the right hand side of the equation in the third case, except ϵ.

Let us turn to the marginal benefit. Up to the first order, the marginal benefit is the increase in the principal's expected net benefits times the probability of new agent types who accept h'. Since the expected payment to the agent of type $\overline{x}$ is $z(\overline{x})$, the increase in the expected net benefits to the principal is $\mathbb{E}[u|\overline{x}] - z(\overline{x})$. To calculate the probability of new agent types who accept h', we need to figure out new types first. Suppose all types between $\overline{x} - \delta x$ and $\overline{x}$ accept h'. That means,

$$z(\overline{x} - \delta x) = (1 + \epsilon)\mathbb{E}[h|\overline{x} - \delta x] - \epsilon z(0).$$

Up to the first order,

$$\delta x = \epsilon \frac{z(\overline{x}) - z(0)}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{d}{dx}z(\overline{x})}.$$

The probability of new agent types accepting h' is $\epsilon \frac{z(\overline{x}) - z(0)}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{d}{dx}z(\overline{x})} p(\overline{x})$, and therefore the marginal benefit is

$$\epsilon \frac{z(\overline{x}) - z(0)}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{d}{dx}z(\overline{x})} p(\overline{x}) \left(\mathbb{E}[u|\overline{x}] - z(\overline{x})\right)$$

This expression is the left hand side of the equation in the third case except ϵ.

The fourth case describes circumstances in which it is optimal for the principal to admit only lower types. Using similar arguments as in the third case, we show that not only the threshold type but also the highest type must be indifferent between accepting and rejecting when the principal minimizes expected payments. The second equation in the fourth case is a first order condition for the threshold type, and has an analogous meaning as the second equation in the third case.

The fifth case is when the principal optimally admits lower types and higher types and rejects middle types. Then, unsurprisingly, the participation constraints of the two threshold types have to be binding. The remaining two equations are the first order conditions for the threshold types.

Two equations can be derived using the argument we used in the third case. The equation (3.1) in the fifth case means that given $\underline{x}$, the principal has no incentive to change $\overline{x}$ via affine transformations. The equation (3.2) in the fifth case implies that given $\overline{x}$, the principal has no incentive to change $\underline{x}$ via affine transformations.

Let us briefly talk about the equation (3.1). Suppose that the principal increases the payment slightly without changing $\underline{x}$. See Figure 3.5. Like before, the left hand side of the equation (3.1) is the marginal benefit to the principal due to new agent types who accept the contract h. The second term in the right hand side of the equation (3.1) is the marginal increase in the payment because the principal pays more to higher types. The first term

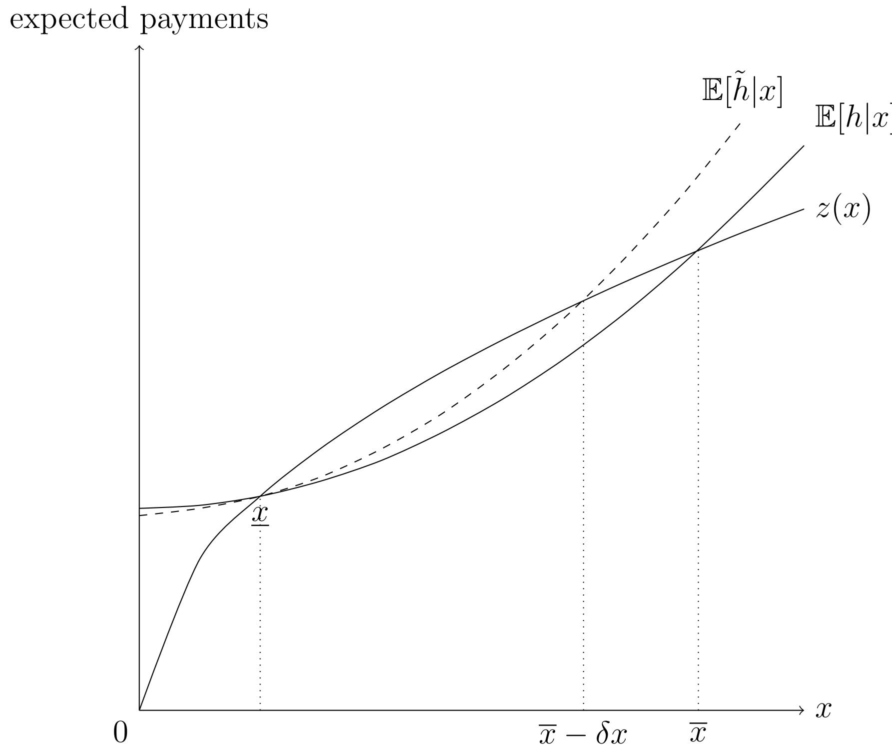

Figure 3.5: Affine transformation that does not alter the lower threshold

in the right hand side of the equation (3.1) is new and corresponds to the marginal reduction of the expected payments. In order to keep $\underline{x}$ unchanged through an affine transformation, the principal necessarily pays less to lower types who accept the contract h. Note that $\mathbb{E}[h|x] \leq z(x)$ for all $x \geq \underline{x}$. Consequently, the right hand side of the equation (3.1) is the net change of the expected costs. Optimality requires that the marginal change of expected benefits and the marginal change of expected payments be the same.

3.7 Linear Contracts

In this section, we consider the special case in which the optimal contract makes the expected payment to the agent a linear function of the agent's type x.

Definition 3.7.1. *A contract h is linear if $\mathbb{E}[h|x]$ is linear in x.*

Lemma 3.7.1. *If at least one incentive compatible contract that is linear exists, then there is an optimal incentive compatible contract that is linear.*

The idea for the proof of Lemma 3.7.1 is shown in Figure 3.6. It shows that, if there is at least one incentive compatible contract, then every incentive compatible contract there is a linear contract that yields at least as high expected benefit to the principal as the original contract. Suppose g is an arbitrary incentive compatible contract. Due to Lemma 3.5.1, $\mathbb{E}[g|x]$ is increasing and convex in x. Now, consider the linear incentive compatible contract the existence of which is assumed in the Lemma 3.7.1. Obviously, there exists an affine transformation of this contract so that it has the same participation thresholds as g. This new contract, say h, is obviously also linear. Moreover, because g is convex, it involves payments that are not higher than the expected payments under g for every agent type who accepts the contract. Therefore, the affine transformation of the given linear contract weakly raises the principal's expected benefits.

The next question is: "When does a linear incentive compatible contract exist?" The following definition helps us to find a sufficient condition under which a linear contract exists.

Definition 3.7.2. *The quality of information is separable if $\xi(x, y) = x\xi_Y(y)$.*

Note that $\xi(x, y) = \xi_X(x)\xi_Y(y)$ can be a functional form such that quality of information is separable. We can redefine $\xi_X(x)$ to be a new agent type,

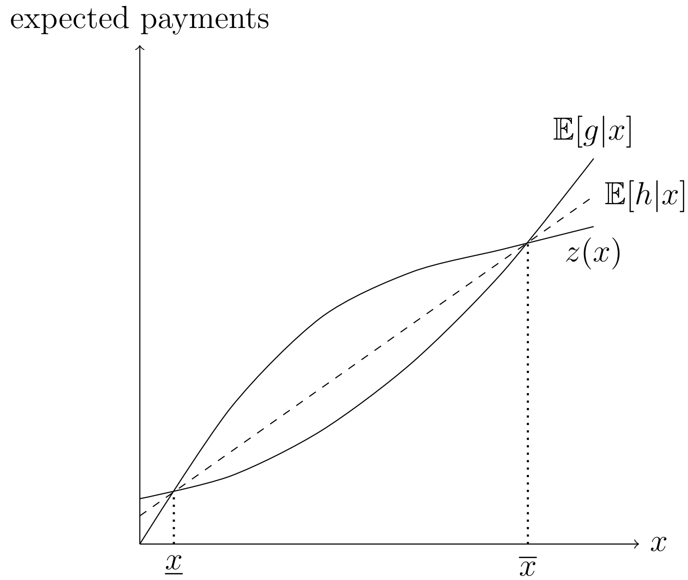

Figure 3.6: Optimal linear contract

and if that re-parametrization does not violate Assumption 1 and Assumption 2, then quality of information becomes separable.

Recall that $\xi(x, y)$ is the Euclidean distance between the prior belief and the posterior belief conditional on (x, y). So, the first order derivative of $\xi(x, y)$ with respect to x, $\frac{\partial}{\partial x}\xi(x, y)$ is the rate at which the posterior belief moves away from the prior belief as x increases. Therefore, the quality of information is separable in the sense of the above definition if and only if the posterior belief moves away from the prior belief at a constant speed as the measure of the quality of information increases. We can now state our result.

Proposition 3.7.1. *If the quality of information is separable, then there exists a linear incentive compatible contract.*

Now that we have established that linear incentive compatible contracts are of special interest in our model, we note the following surprising feature of linear contracts. The following proposition holds regardless of whether the linear incentive compatible contract considered in the result is optimal or not.

Proposition 3.7.2. *If h is a linear incentive compatible contract, and if S is a corresponding proper scoring rule, then S is discontinuous at the prior μ_0, that is, there is at least one sequence of posteriors $(\vec{\mu}_n)_{n \in \mathbb{N}}$ that converges to the prior belief $\vec{\mu}_0$, and a state $\omega \in \Omega$ such that:*

$$\lim_{n \to \infty} S(\omega, \vec{\mu}_n) \neq S(\omega, \vec{\mu}_0).$$

The proof of Proposition 3.7.2 is indirect. The proof shows that, if the scoring rule were continuous at the prior belief, then the agent's expected payoff, as a function of the quality x of the agent's information, would have slope zero at $x = 0$. But for a linear contract, this slope has to be strictly positive. Therefore, we obtain a contradiction. The reason why continuity at the prior implies zero slope of expected payments at at $x = 0$ is that continuity of the scoring rule would allow us to apply the envelope theorem. Because reporting $\vec{\mu}_0$ is optimal at $x = 0$, when exploring the effect of a change in x on the expected payoff, the effects of the change in reporting strategy are locally of second order. We might as well assume that the agent continues to report the prior even if x increases slightly. But if the agent continues to report the prior, then his expected payment is not going to change.

One can view the agent's decision problem in our model as a special case of the model in Radner and Stiglitz (1984). The agent in our model corresponds to a decision maker in Radner and Stiglitz (1984). The proper scoring rule indicates the agent's utility. The agent's action set consists of all possible posterior beliefs. The agent's type corresponds to the parameter that indexes a family of information structure in Radner and Stiglitz (1984).[8] Radner and Stiglitz's famous Theorem (Radner and Stiglitz (1984, p. 36)) says that in their model, the marginal value of information is zero if we start with a totally uninformative information structure. The proof of their result is based on the envelope theorem argument that we sketched in the previous paragraph. Their result cannot apply to our model because, in a linear contract, the marginal value of information at $x = 0$ is positive. Therefore, one of Radner and Stiglitz's assumptions must be violated. Proposition 3.7.2 identifies the assumption that is violated.[9]

Proposition 3.7.2 is phrased in terms of scoring rules, but it is also interesting to understand why Radner and Stiglitz's theorem cannot be applied to the agent's decision problem when we adopt the alternative perspective where the pair (x, y) is the agent's choice variable, and the agent's payoff is determined by the contract h directly. It turns out that, from this perspective, the assumptions of Radner and Stiglitz's theorem are violated for *every* incentive compatible contract, not just for linear incentive compatible contracts. This is a somewhat trivial observation that follows from the definition of incentive compatibility. We explain the argument in the next paragraph. But first, let us clarify the logical connection between the observations just made. Radner and Stiglitz's theorem provides a *sufficient* condition for the marginal value of information to be zero. For linear contracts, we know that this sufficient condition must be violated, because the conclusion of the theorem is violated. This is not interesting if we treat the pair (x, y) as the agent's choice variable, because in that case the sufficient condition is *always*

[8] Radner and Stiglitz's parameter θ.

[9] In Radner and Stiglitz's notation: they assume that for every state s the function $u_s(\cdot, \cdot)$ is continuous on the Cartesian product of action and type space. In our setting, u_s corresponds to the scoring rule. It does not depend on the agent's type, and is therefore trivially continuous in its second argument. Because our model satisfies all other assumptions of Radner and Stiglitz, for some s continuity of u_s in the first argument must be violated.

violated, regardless of whether the contract is linear or not. By contrast, if we consider the scoring rule, the sufficient condition is not *always* violated, and therefore the conclusion of Proposition 3.7.2 is non-trivial.

It remains to explain why, for incentive compatible contracts h, if we treat the pair (x, y) as the decision variable, the assumptions of Radner and Stiglitz's theorem are automatically violated, regardless of whether the contract is linear or not. The reason is that Radner and Stiglitz's theorem assumes that the agent's choice function, that maps agent type and observed signal into an action, is continuous in both arguments, and "flat" at $x = 0$. In our case, for an incentive compatible contract, the choice function is the identity mapping. Of course, it is continuous, but it is not "flat" at $x = 0$. Radner and Stiglitz mean by "flat" that the choice is independent of the signal. But, for an incentive compatible contract, even if $x = 0$, the agent reports the signal y truthfully, and therefore the decision function is not flat. One might see this as an artificiality. We might as well assume that, when the agent tells the principal that he does not have an informative signal, i.e. reports $x = 0$, he reports the same signal realization regardless of which realization he has actually observed. But, if we make that assumption, the decision function is not continuous, because, for $x > 0$, it matters that the agent does report the signal correctly.

We conclude by mentioning an interesting implication of Proposition 3.7.2. In Radner and Stiglitz's environment, an uninformed decision maker has no incentive to buy a small piece of information if the information is priced linearly in the quality of that information. However, if the contract is linear, an uninformed agent will have an incentive to buy even a small piece of information, provided that the price of the quality of information is low enough.

3.8 Example

We present a simple example and analyze it. There are two states, ω_1 and ω_2. The set of agent types is the same as before: $X = [0, 1]$. The agent's signal has only two realizations: y_1 and y_2. The agent's outside option is $z(x) = \frac{1}{2}\sqrt{x}$. The table below shows the probabilities of receiving a signal realization conditional on the state and the agent type. Initially, the two states are equally probable.

	ω_1	ω_2
y_1	$\frac{1}{2}(1 + x)$	$\frac{1}{2}(1 - x)$
y_2	$\frac{1}{2}(1 - x)$	$\frac{1}{2}(1 + x)$

Conditional on (x, y_1),

$$\begin{aligned}
\mu(\omega_1|x, y_1) &= \frac{1}{2}(1 + x), \\
\mu(\omega_2|x, y_1) &= \frac{1}{2}(1 - x).
\end{aligned}$$

Conditional on (x, y_2),

$$\begin{aligned}
\mu(\omega_1|x, y_2) &= \frac{1}{2}(1 - x), \\
\mu(\omega_2|x, y_2) &= \frac{1}{2}(1 + x).
\end{aligned}$$

The principal's utility is given below:

$$\begin{aligned}
u(\omega_1, a) &= -4a^2 + 1, \\
u(\omega_2, a) &= -4(1 - a)^2 + 1,
\end{aligned}$$

where $a \in [0, 1]$. Given the probability q that the true state is ω_1, the principal's optimal action $a^*(q) = 1 - q$. Therefore, if the principal chooses optimally given q, $\mathbb{E}[u|q] = -4q(1-q)+1$. The principal's utility is normalized such that at the prior $q = 0.5$, $\mathbb{E}[u|q = 0.5] = 0$. We need to describe the principal's expected utility in terms of x. With probability 0.5, y_1 is observed, and the posterior belief is $\frac{1}{2}(1+x)$. With probability 0.5, y_2 is observed, and the posterior belief is $\frac{1}{2}(1-x)$. Substituting these posteriors into the formula for $\mathbb{E}[u|q]$ and taking expected values, we obtain that the principal's expected utility, conditional on x, is: $\mathbb{E}[u|x] = x^2$.

It is easy to check that the posterior beliefs satisfy Definition 3.7.2, i.e., the quality of the information is separable. So, Proposition 3.7.1 implies that a linear incentive compatible contract exists, and Lemma 3.7.1 means that an optimal contract can be obtained through an affine transformation of a linear incentive compatible contract.

We characterize an optimal contract using the following contract g.

g	ω_1	ω_2
y_1	1	-1
y_2	-1	1

First, we show that g is incentive compatible. Consider (x, y_1). Note that we do not have to worry about cases in which the agent reports (x', y_1) because g is independent of x. When the agent truthfully reports (x, y_1), his expected utility is

$$
\begin{aligned}
& g(\omega_1, x, y_1)\mu(\omega_1|x, y_1) + g(\omega_2, x, y_1)\mu(\omega_2|x, y_1) \\
&= \frac{1}{2}(1+x) - \frac{1}{2}(1-x) \\
&= x.
\end{aligned}
$$

If the agent reports (x', y_2),

$$
\begin{aligned}
& g(\omega_1, x', y_2)\mu(\omega_1|x, y_1) + g(\omega_2, x', y_2)\mu(\omega_2|x, y_1) \\
&= -\frac{1}{2}(1+x) + \frac{1}{2}(1-x) \\
&= -x.
\end{aligned}
$$

So, the agent has a weak incentive to report (x, y_1). Due to the symmetry of the example, the agent also has an incentive to truthfully report when his type is x and observes y_2. Therefore, g is incentive compatible.

Second, we show that under g, the expected payment to the agent is linear in x. The above calculation shows that the expected payment to the agent of type x is x. That is, $\mathbb{E}[g|x] = x$.

Third, we numerically solve for the two optimal thresholds. The numerical solution tells us that $\underline{x} \approx 0.017$ and $\overline{x} \approx 0.706$. The corresponding optimal contract, that is, the optimal affine transformation of g, is shown in the following table.

$\tilde{g}$	ω_1	ω_2
y_1	1.08	0.05
y_2	0.05	1.08

Figure 3.7 shows the principal's expected utility, the expected payment to the agent, and the agent's outside option as functions of x.

Interestingly, Figure 3.7 implies that the principal admits lower types and pays the agent more than his outside option. In fact, the principal actually loses in case the agent turns out to be a bad forecaster. Also, note that the principal does not hire types between $[x^\dagger, \overline{x}]$, for which she is able to pay the agent more than his outside option and benefit from the agent's information. This is because the principal wants to reduce the expected payment to higher

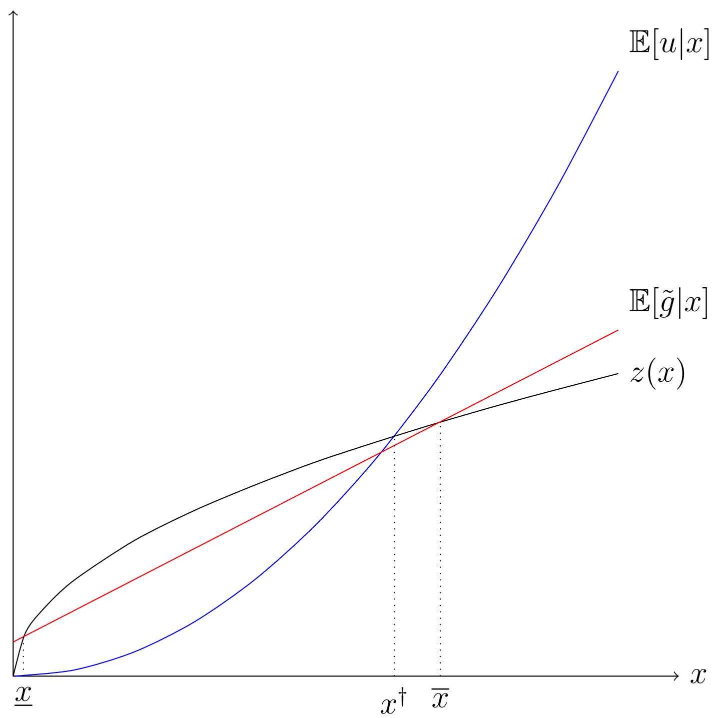

Figure 3.7: Comparison of the principal's expected utility, the expected payment to the agent and the agent's outside option: $\tilde{g}$ is the best affine transformation of g.

types. By rejecting mediocre types and admitting extremely low types, the principal can reduce the payment to higher types as well as ex-ante expected costs. In this example, the reduction in overall costs outweighs the loss in the utility of the principal. Figure 3.8 shows the expected net benefits to the principal.

We show the shape of a proper scoring rule S that corresponds to g. Since there are only two states, a posterior belief can be represented by $q \in [0, 1]$, probability of ω_1. That is, a proper scoring rule can be a function from $\Omega \times [0, 1]$ to the extended real line. We use this definition to draw S. Note that when the agent receives y_1, the probability q of state ω_1 is larger than

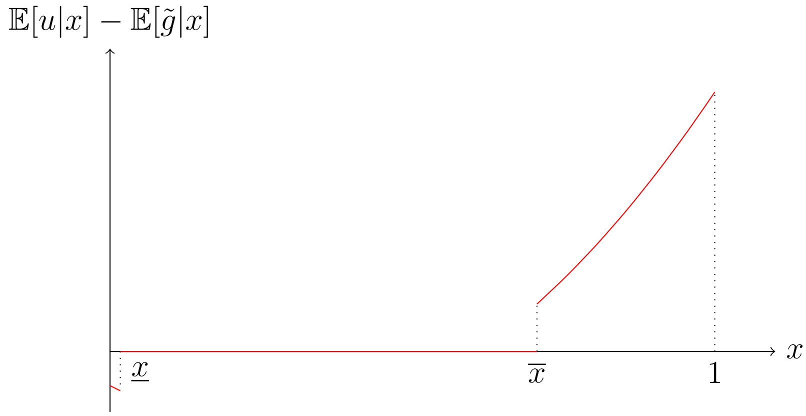

Figure 3.8: Expected net benefits to the principal under $\tilde{g}$

0.5, and he receives 1 if ω_1 materializes. If the agent observes y_2, q is less than 0.5. If ω_1 is realized, the agent pays 1 to the principal. If the agent is uninformed, he is indifferent between -1 and 1. This is depicted in Figure 3.9. One can obtain the graph of $S(\omega_2, \cdot)$ by reflecting $S(\omega_1, \cdot)$ across the horizontal axis. As Proposition 3.7.2 suggests, $S(\omega, q)$ is discontinuous at $q = 0.5$ for all $\omega \in \Omega$.

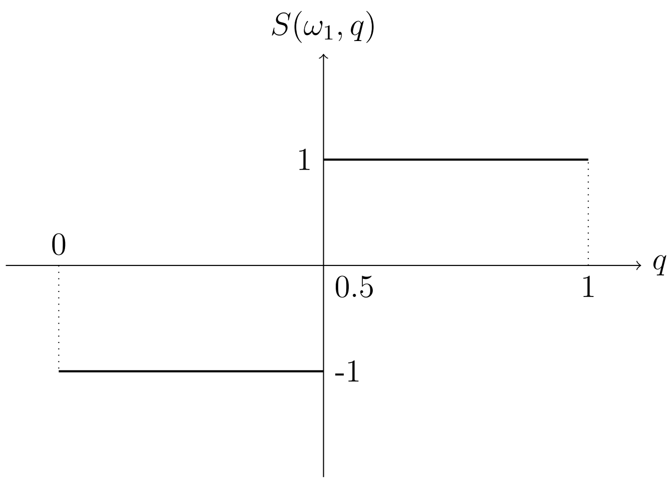

Figure 3.9: Proper scoring rule that corresponds to g

We close this section, providing a complete characterization of proper scoring rules in this example. The theorem due to McCarthy in Section 3 *indirectly* means that for any convex function $\phi : [0,1] \longrightarrow \mathbb{R}$ and a subgradient ϕ' of ϕ, the following scoring rule S is proper:[10]

$$S(\omega_1, q) = \phi(q) + (1-q)\phi'(q),$$
$$S(\omega_2, q) = \phi(q) - q\phi'(q).$$

Note that this characterization works for all decision problems with two states and two signal realizations. Information structure does not matter. It is possible that a convex function ϕ is not differentiable at the prior. In this case, the right limit of $\phi'(q)$ and the left limit of $\phi'(q)$ are different at the prior, and the above two equations mean that the proper scoring rule derived from ϕ is discontinuous at the prior.

3.9 Discussion of Assumptions 1 and 2

In this section, we discuss Assumptions 1 and 2. At first glance, they look restrictive. An easy way of relaxing Assumption 1 is to discard it, i.e., to impose no assumption on the agent's type. However, if there is no structure on the agent's type, the agent's type cannot be interpreted as the quality of information. In this case, it would be difficult to interpret results and understand intuitions.

Blackwell dominance would be a good candidate that can replace Assumption 1. That is, a higher type Blackwell-dominates a lower type. By imposing Blackwell dominance, the agent's type can be thought of as the quality of information. Also, Blackwell dominance is a weaker assumption in the sense that it implies that the expected utility increases in the agent's

[10] According to the original definition we use in Section 3, ϕ takes two arguments. This is why we use the word "indirectly".

type whereas Assumption 1 implies that the expected utility increases in the agent's type at increasing rate (Lemma 3.5.1). However, without further assumptions, it would be more complicated to analyze the agent's participation constraint. Even if we maintain Assumption 2, increasing expected utility does not necessarily mean that there are at most two thresholds. If we only assume Blackwell dominance, there could be more than two thresholds, and characterizing thresholds would require a mathematically cumbersome calculation of the model.

Another possibility is to assume Blackwell dominance and other assumptions so that the marginal benefit of higher quality of information is decreasing. In other words, assumptions can be made such that the expected payment to the agent is increasing in the agent's type at decreasing rate. Also, instead of Assumption 2, one can assume that the agent's outside option is convex in the agent's type. This corresponds to the increasing marginal cost of producing information. If we imagine the typical behavior of marginal benefits and marginal costs, those assumptions are reasonable. However, we need to address one issue, which is that it is not easy to make the marginal benefit of information increasing everywhere. This is because of Radner and Stiglitz's result, which implies that the marginal value of information is zero at zero type. So, if we want to work with conventional marginal benefits and marginal costs, we need to resolve the issue. Especially, we need to find an "inflection" type, before which the marginal value of information is increasing and after which the marginal value of information is decreasing. In case we can successfully characterize such an inflection type, we would be able to characterize the agent's participation constraint.

With the assumptions made in the previous paragraph, it is possible that the principal hires intermediate types only. For instance, an inflection type can be very small. In this case, our analysis remains valid. When the quality of information is separable, an optimal contract is a linear contract. The existence of a linear contract is independent of Assumptions 1 and 2. Figure 3.10 clearly shows that if a linear contract exists, the optimal contract is a linear contract. In Figure 3.10, z is the agent's outside option, g is a linear contract, and h is an incentive compatible contract. Only types between $\underline{x}$ and $\overline{x}$ accept h. g is obviously better than h because any type who accepts h also accepts g and the principal pays less to the agent conditional on the acceptance.

Assuming that an inflection type is small enough, we describe how Proposition 3.6.1 needs to be changed. If it is optimal for the principal to admit no types, the first part in Proposition 3.6.1 remains true. If it is optimal for the principal to hire all types, the participation constraint for the lowest type and the highest type should be binding. Otherwise, the principal can lower transfers to the agent until the participation constraints for the lowest and the highest types are binding. For the third case, the participation constraints for the threshold type and the highest types are binding. For the fourth case, the participation constraints for the threshold type and the lowest types are binding. For the fifth case, the two threshold types are indifferent between accepting and rejecting. Regarding the equations that thresholds satisfy, one can derive new equations using the new participation constraint and the equalization of marginal benefit and marginal cost.

3.10 Conclusion

In this section, we summarize our finding and its implications. Our result characterizes optimal contracts for the principal under the following two assumptions. The marginal benefit of higher quality of information is increasing and the marginal cost of producing higher quality of information is decreasing.

First, we characterize a relationship between incentive compatible con-

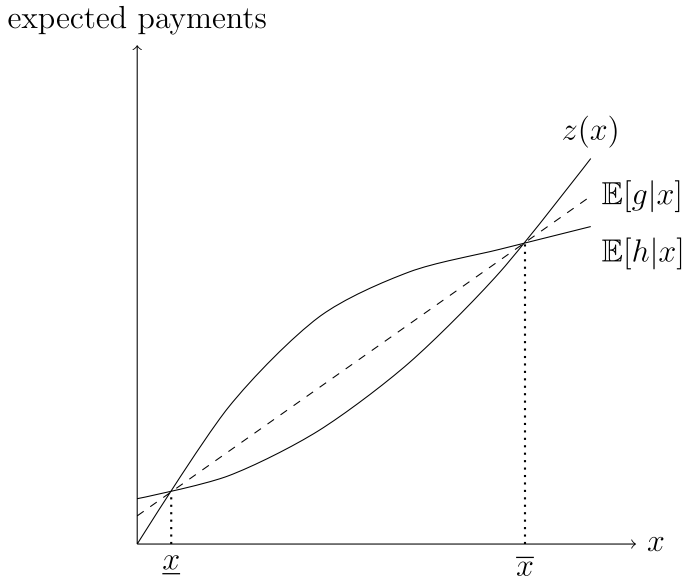

Figure 3.10: Optimal linear contract under different assumptions

tracts and proper scoring rules. We show that given an incentive compatible contract, one can find a corresponding proper scoring rule and that given a proper scoring rule, one can find an incentive compatible contract.

Second, we provide a necessary condition for an optimal contract. Under our assumptions, there are at most two threshold types ($\underline{x}$ and $\overline{x}$) such that it is optimal for the principal to not hire types between the two threshold types. Since the agent type is related to the quality of information, the principal hires bad forecasters and good forecasters, but are unwilling to hire mediocre forecasters.

Third, we show that if the quality of information is separable, a linear contract exists and an optimal contract can be obtained through an affine transformation of a linear contract. The existence of a linear contract means the violation of Radner and Stiglitz's result (1984).

By definition of a linear contract, the marginal value of information is positive when a decision maker is uninformed. The reason that the marginal value of information is positive is because the assumption that a decision function is continuous at the prior is violated. One can show that a proper scoring rule that corresponds to a linear contract must be discontinuous at the prior. In a Radner and Stiglitz setting, the decision maker is not willing to buy an infinitesimal amount of information, regardless of the information price. However, in our model, if a linear contract exists, the marginal value of information is positive. As a result, the decision maker would be willing to buy a small amount of information provided that the price is not too high.

In this work, we make a special assumption on the quality of information to simplify the analysis of the agent's participation constraint. In order to make our model more convincing, one necessary extension is to relax Assumption 1 and come up with a less restrictive assumption on the quality of information. Also, it would be important to find sufficient conditions for the optimal contract. Another extension of this paper is to consider a risk averse agent. If the agent is risk averse, proper scoring rules do not induce the agent to truthfully report, which means that a characterization of incentive compatible contracts would be more difficult. In addition, the agent's risk preference may have some impact on the participation constraint. So, one future project is to investigate the impact of the agent's risk preference on the optimal contract.

Chapter 4

On The Optimal Order of Experiments

This chapter is joint work with Professor Tilman Börgers.

4.1 Introduction

Consider a decision maker who wants to make a decision under uncertainty, but, before choosing, has the option of observing the realization of two signals. Observing the realization of a signal is costly. Thus, if it is optimal to observe both signals, it may be a cost saving strategy for the decision maker to observe one signal first, and to make the decision whether or not to observe the second signal depending on the realization of the first signal. In which order should the decision maker observe the two signals? For example, if one signal is unambiguously more informative than the other, should the decision maker observe the more informative signal first?

At first sight the answer to this question seems to necessarily be: "it depends." It seems unlikely that any results can be obtained that don't depend on the the details of the distributions of the two signals, as well as the nature of the decision problem. A formal statement of the claim that nothing much can be said in general is a theorem due to Greenshtein. Greenshtein used the terminology of the statistics literature and referred to signals as "experiments." Greenshtein (1996, Theorem 3.2) then says that for two independent experiments F and G one can say that F should be run before G regardless of what the underlying decision problem is if and only if F is essentially equivalent to running G and some other experiment that is independent of G. Thus, cases in which the order in which F and G should be conducted is independent of the underlying decision problem are very rare.

Greenshtein illustrates his result with two examples. In the first example both F and G are binary. Then nothing can be said about the optimal order of F and G independent of the decision problem (except in the trivial case in which $F = G$, in which the order does not matter). The second example is the case in which F and G are normally distributed. In this case, Greenshtein's theorem implies that the decision maker always wants to run the experiment with the smaller variance first.

In this paper, we revisit Greenshtein's result and investigate whether more can be said about the optimal order of experiments when the class of decision problems that is considered is not quite as large as in Greenshtein's paper. Specifically, Greenshtein allowed decision problems in which the decision maker takes actions twice, after observing the first signal, and after observing the second signal, and action spaces and payoff functions were essentially unrestricted. In our paper, we assume that the information gathering process precedes any decision making. Moreover, we restrict attention to the case that the underlying state space as well as the set of available actions are binary. The leading example that we have in mind is the "jury problem" that has been much studied in economics: a defendant is either guilty or not guilty, and the jury has to find the defendant either guilty or not guilty.

If the class of decision problems is restricted, one can hope to find that the optimal order in which to conduct experiments can be determined more often than when the class of decision problems is unrestricted. In this paper we investigate this in detail. For the case of binary signals we find that Greenshtein's result continues to be true. We offer a detailed investigation that clarifies the intuition for this result. If a signal has more than two realizations, by contrast, Greenshtein's result is no longer true. We display a class of experiments for which we can determine the optimal order in which experiments can be conducted regardless of the nature of the underlying decision problem, as long as it belongs to the class of decision problems to which we have restricted our attention.

The results described so far assume that each experiment has some fixed cost. Moreover, in order to make the problem interesting, we assume that the cost of the two experiments are identical, so that the optimal order of experiments is not driven by cost differences. In a second part to the paper, we investigate the case in which the cost of information gathering is due to discounting, that is, the more information the decision maker gathers, the longer it takes until a decision is made, and postponing the decision is costly. We show that when discounting is the cost of waiting, the basic structure of the results remains the same as in the cost of fixed cost, except that we can obtain slightly stronger positive results, that is, results that permit decision problem independent results on the optimal order of experiments.

One may view the relation between this paper and Greenshtein (1996) as analogous to the relation between Athey and Levin's (2018) analysis of the informativeness of signals in a specific class of decision problems, and Blackwell's (1953) paper.[11] In fact, there is a close formal connection between Greenshtein's and Blackwell's result, although there is no formal connection between our and Athey and Levin's results.

In addition to revisiting Greenshtein's result in a more restricted class of decision problems, we also go beyond his analysis by providing a detailed

[11] On this relation, see Greenshtein (1996) and also Oliveira (2018).

analysis that illuminates how negative or positive results arise. To explain the intuition for our results, let us focus first for the moment on the case that experiments have only two realizations. Suppose that the two experiments are unambiguously ordered in terms of (static) informativeness: one of the experiments Blackwell dominates the other. It is easy to provide examples of decision problems in which it is optimal to run the more informative experiment first. The intuition is that this experiment often settles the question of what to do, and therefore no cost of running a second experiment needs to be incurred. The more interesting point is that we show that even if the class of decision problems is restricted, there are cases in which the decision maker prefers running the less informative experiment first. This implies that no statement can be made that is independent of the underlying decision problem.

To understand the intuition why one might run the less informative experiment first, let us think of a venture capitalist who is about to decide whether to invest seeding funds into a startup company. The venture capitalist believes that the startup company will be successful with low probability, but conditional on the success, it will generate large income flows. The venture capitalist has two sources of information: her friend and a professional analyst. Assume that the analyst is more informative than the venture capitalist's friend. Suppose that if the venture capitalist collects no information, her optimal action is to invest, and that the venture capitalist does not invest only if her friend and the analyst recommend her to not invest. In other words, if either the friend or the analyst recommends the investment, the venture capitalist's optimal choice is to invest seeding funds into the startup company. Given the venture capitalist's belief, if she asks the analyst about what to do first, the analyst recommends to her to not invest with high probability, which means that with high probability, the venture capitalist also asks her friend. However, if the venture capitalist talks to her friend first, her friend recommends her to invest with high probability. As a result, there is only a small chance that after talking to her friend, the venture capitalist also visits the analyst. In this circumstance, if the venture capitalist incurs fixed costs, she can save costs by asking her friend first. If discounting is the cost of experimentation, the venture capitalist can take the right action earlier by talking to her friend first.

The above situation illustrates the following idea: in some cases, a more informative experiment triggers the second experimentation more often than a less informative experiment does. Therefore, one might find it desirable to run the more informative experiment first. It turns out that in the case that experiments have only two realizations, it is always the case that the problem described in the previous example arises, and therefore the decision maker wants to run the less informative experiment first. But this is somewhat artificial. If experiments have more than two realizations, then it is possible to construct classes of examples in which it is unambiguously better to run the more informative experiment first. This is the intuitive basis of our second result.

The paper is organized as follows. We present our model in Section 2. In Section 3, we discuss a normalization of payoffs that will simplify the later derivation of results, and argue that fixed costs and discounting have different impacts on the decision maker's incentive to run an experiment. In Section 4, we investigate in detail a particular case of binary experiments when the decision maker incurs fixed costs of experimentation and does not discount future payoffs. The purpose of this section is to develop intuition for the optimal order of experiments. In Section 5, we prove that, if experiments have fixed costs, then Greenshtein's Theorem 3.2 remains valid in our set-up. In Section 6, we provide a class of examples in which one signal has three rather than two realizations, and in which Greenshtein's Theorem 3.2 is no longer valid. In Sections 7, 8 and 9, we consider the case that the cost of experimentation is due to discounting, and not to fixed costs. Because of the difference between fixed costs and discounting, new calculations are required to analyze the optimal order of experiments when discounting is the only cost of experimentation. In Section 9, we obtain slightly more positive results. Section 10 is a conclusion.

4.1.1 Related Literature

Our model is related to dynamic information acquisition. DeGroot (1962) studies the optimal sequence of experimentation, and showed if one experiment Blackwell-dominates all other experiments, the optimal sequence is to repeat the "best" experiment in each period. Aghion *et al.* (1991) considers a decision maker who is uncertain about her payoff function and investigate whether she can asymptotically learn her payoff function. Che and Mierendorff (2017) and Mayskaya (2018) consider a situation in which a decision maker with a limited amount of attention has several sources of information and characterize optimal policies that assign the amount of attention to each source of information at every point of time. Liang *et al.* (2018) study Gaussian experiments with arbitrary correlation and characterize conditions under which "myopic" information acquisition is optimal. In these papers, the decision maker can run the same experiment in every period. In our model, the decision maker cannot run the same experiment unless the set of available experiments contains two identical experiments. In other words, the set of available experiments becomes smaller whenever the decision maker runs an experiment.

This paper is also related to the comparison of experiments. Blackwell (1953) proposes the notion of informativeness in static decision problems and provides the partial ordering among experiments. Blackwell's notion can be naturally extended to the comparison of deterministic sequences of experiments, and Greenshtein (1996) and Oliveira (2018) show that one sequence of experiments is better than another sequence of experiments regardless of decision problems if and only if the second sequence of experiments can be sequentially reproduced from the first sequence of experiments.[12] In this paper, we consider a class of decision problems, and as a result, Greenshtein's results (Theorem 1.1a and Theorem 3.2, 1996) do not have to hold. We find a class of experiments that violates Greenshtein Theorem 3.2, and for which there exists a non-trivial comparison of experiments.

There are several papers that extend the Blackwell's partial order of experiments by considering a class of decision problems. Lehmann (1988) studies monotone decision problems and obtains an informativeness ordering for the class of monotone decision problems. Athey and Levin (2018) extends Lehmann's work. Cabrales, Gossner and Serrano (2013) considers an agent who has weakly increasing relative risk aversion, and shows that the value of an experiment is equivalent to the reduction in entropy due to the experiment. This means that they obtain the complete ordering because reduction in entropy is a number. Ganuza and Penalva (2010) introduce two definitions of informativeness that are different from Blackwell's definition of informativeness. Our model shares a similar spirit with those papers in the sense that we restrict our attention to a class of decision problems.

4.2 Model

Environment Nature chooses a state $\omega \in \Omega = \{A, B\}$ at the outset. This choice is not revealed to the decision maker, but she knows $\mu_0 \in (0, 1)$, the initial probability that the true state is A. The decision maker can only run at most two experiments, but she can run one experiment at a time.

In period 0, she chooses between taking an action and running an experiment. Once the decision maker takes an action, the decision process is over.

[12] Oliveira (2018) proves Greenshtein (1996) Theorem 1.1a using a different approach.

The true state is revealed and the decision maker receives utility based on her action and the true state. If the decision maker runs an experiment, she moves to period 1. The decision maker observes the experiment outcome, and chooses between an action and the remaining experiment. If the decision maker runs the remaining experiment, she moves to period 2, observes a realization, and chooses an action.

The decision maker incurs fixed costs c whenever she conducts an experiment. Future payoffs are discounted by δ.

Experiments An experiment F is represented by 3-tuple (X, f_A, f_B). X is a finite set of outcomes. For each $\omega \in \Omega$, f_ω is a probability measure on sigma field of X. Two experiments F and G are available, and they are independent conditional on the state. We assume that for all $\omega \in \Omega$, all $x \in X$ and $y \in Y$, $f_\omega(x)g_\omega(y) > 0$.

Binary Decision Problems The decision maker can choose an action $\alpha \in \mathcal{A} = \{a, b\}$. Her payoff matrix u can be represented by 2×2 matrix:

u	A	B
a	$u_A(a)$	$u_B(a)$
b	$u_A(b)$	$u_B(b)$

We assume that $u_A(a) > u_A(b)$ and $u_B(a) < u_B(b)$ so that the decision maker wants to choose the action to match the state. Also, we assume that $u_A(a)$ and $u_B(b)$ are positive. A binary decision problem $\mathcal{D}$ is the tuple (u, c, δ). Note that a binary decision problem does not include the decision maker's prior belief.

A Sequence of Experiments We use FG to denote a sequence of experiments. FG means that if the decision maker runs an experiment, she runs F first. Similarly, GF means that, if she runs an experiment, she runs G first.

The Set of Utility Vectors We construct the set of utility vectors given the sequence of experiments FG. Let σ^0_{FG} be the decision maker's strategy in period 0:

$$\sigma^0_{FG} : \mathcal{A} \cup \{F\} \longrightarrow [0,1]$$

where $\sum_{\alpha \in \mathcal{A} \cup \{F\}} \sigma^0_{FG}(\alpha) = 1$. Σ^0_{FG} is the set of all strategies σ^0_{FG} in period 0.

Let σ^1_{FG} be the decision maker's strategy in period 1:

$$\sigma^1_{FG} : (\mathcal{A} \cup \{G\}) \times X \longrightarrow [0,1]$$

where for each $x \in X$, $\sum_{\alpha \in \mathcal{A} \cup \{G\}} \sigma^1_{FG}(\alpha, x) = 1$. That is, $\sigma^1_{FG}(\cdot, x)$ assigns probability to the action in $\mathcal{A}$ or G given the outcome x of F. Σ^1_{FG} is the set of all strategies σ^1_{FG} in period 1.

σ^2_{FG} is a strategy in period 2:

$$\sigma^2_{FG} : \mathcal{A} \times X \times Y \longrightarrow [0,1]$$

where for each $(x, y) \in X \times Y$, $\sum_{\alpha \in \mathcal{A}} \sigma^2_{FG}(\alpha, x, y) = 1$. Let Σ^2_{FG} be the set of all strategies σ^2_{FG} in period 2, and $\Sigma_{FG} = \Sigma^0_{FG} \times \Sigma^1_{FG} \times \Sigma^2_{FG}$.

Given the strategy $\sigma_{FG} \in \Sigma_{FG}$, we define a utility vector. For each state $\omega \in \Omega$,

$$\begin{aligned} v_\omega(\sigma_{FG}) &= \sum_{\alpha \in \mathcal{A}} u_\omega(\alpha)\sigma^0_{FG}(\alpha) + \sigma^0_{FG}(F)\left(-c + \delta \sum_{\substack{\alpha \in \mathcal{A}, \\ x \in X}} u_\omega(\alpha)\sigma^1_{FG}(\alpha, x) f_\omega(x)\right) \\ &+ \sigma^0_{FG}(F)\delta \sum_{x \in X} \sigma^1_{FG}(G, x) f_\omega(x) \left(-c + \delta \sum_{\substack{\alpha \in \mathcal{A}, \\ y \in Y}} u_\omega(\alpha)\sigma^2_{FG}(\alpha, x, y) g_\omega(y)\right) \end{aligned}$$

This is the ex-ante expected utility when the true state is ω and the strategy σ_{FG} is chosen.

Given the strategies σ_{FG}, the utility vector is

$$v(\sigma_{FG}) = (v_A(\sigma_{FG}), v_B(\sigma_{FG})).$$

Given a binary decision problem $\mathcal{D} = (u, c, \delta)$, the set of utility vectors for FG is denoted by $\mathcal{V}_{FG}$.

$$\mathcal{V}_{FG}(\mathcal{D}) = \{v(\sigma_{FG}) | \sigma_{FG} \in \Sigma_{FG}\}.$$

Now, we define a notion of informativeness for a sequence of experiments in our model. Since we are going to consider a subclass of binary decision problems and look for the comparison of sequential experiments, the following definition is tailored to the class of binary decision problems.

Definition 4.2.1. *Let $\mathscr{D}$ be a set of binary decision problems. FG is more informative than GF with respect to $\mathscr{D}$ if $\mathcal{V}_{GF}(\mathcal{D}) \subseteq \mathcal{V}_{FG}(\mathcal{D})$ for all $\mathcal{D} \in \mathscr{D}$.*

This definition is the same definition of informativeness as in Blackwell (1953). Therefore, the following statement is true due to Blackwell's result. FG is more informative than GF with respect to $\mathscr{D}$ if and only if the optimal strategy for the decision maker is either to run F first or to take the right action without experimentation in every binary decision problem in $\mathscr{D}$.

4.3 Normalization of Payoffs and the Cost of Experimentation

In this section, we first show that if fixed costs are the only cost of conducting an experiment, we can restrict our attention to a simple subclass of binary decision problems. This will be done by normalizing the payoff matrix. Then, we discuss why we cannot focus on the subclass of binary decision problems when the decision maker discounts future payoffs.

To normalize the payoff matrix, we need to understand optimal strategies. In our model, the optimal strategy is simple. If the benefit of running an experiment is greater than cost of running the experiment, then the decision maker chooses to run the experiment. Otherwise, the decision maker choose the right action. Since there are only finite periods, we can find an optimal strategy by applying the argument backwards.

To state the result, we need notations. u is the same as in Model section. For given u, define β:

$$\beta = \frac{u_B(b) - u_B(a)}{u_A(a) - u_A(b)}$$

and a new payoff matrix u_β:

u_β	A	B
a	1	0
b	0	β

.

Let $\tilde{c} = \frac{c}{u_A(a) - u_A(b)}$. Given $\mathcal{D} = (u, c, 1)$, let $\tilde{\mathcal{D}} = (u_\beta, \tilde{c}, 1)$. Define $\mathscr{D}_{(u,c)} = \{(u, c, 1)\}$. We denote by $\sigma^*(\mu, \mathcal{D})$ the optimal strategy for the decision maker given her prior μ_0 and the binary decision problem $\mathcal{D}$.

Lemma 4.3.1. *For every prior μ and every binary decision problem $\mathcal{D} \in \mathscr{D}_{(u,c)}$, $\sigma^*(\mu, \mathcal{D}) = \sigma^*(\mu, \tilde{\mathcal{D}})$.*

(proof) In the Appendix.

Lemma 4.3.1 is a special type of normalization of payoff matrices. To understand Lemma 4.3.1, we need to understand the optimal strategy for the decision maker. An optimal strategy is simple. If it is not beneficial for the decision maker to run an experiment, the decision maker chooses the right action as a function of the posterior belief. If running an experiment is beneficial, the decision maker should run the experiment. Lemma 4.3.1 means that such a normalization does not change both of the decision maker's right action and the incentive to run an experiment.

Let us first look at the right action. Let

$$\mu^* = \frac{\beta}{1+\beta}.$$

In both binary decision problems $\mathcal{D}$ and $\tilde{\mathcal{D}}$, the decision maker's right action depends only on whether her posterior belief is larger than or less than μ^*. Assuming the decision maker chooses a when she is indifferent between the two actions a and b, if the posterior belief is larger than or equal to μ^*, the right action is a. If the posterior belief is less than μ^*, the decision maker is supposed to choose b.

Now, let us move onto the decision maker's incentive to run an experiment. In order to normalize u to u_β, we need to subtract the state-dependent payoff v from u and multiply $u - v$ by a positive number. That is, for each $\omega \in \Omega$,

$$u_{\beta,\omega}(a) = \gamma(u_\omega(a) - v_\omega).$$

Since v is independent of the action and the signal realization, the impact of experimentation on u_β is only through u. This means that if an experiment increases the expected utility by 1 in the original binary decision problem $\mathcal{D}$, the experiment increases the expected utility by γ in the new binary decision problem $\tilde{\mathcal{D}}$. Therefore, if we multiply c by γ, the incentive for the principal to run an experiment is preserved under the normalization.

Note that the argument is not unilateral. We can construct u from u_β. So, the optimal strategy in the binary decision problem $\mathcal{D}$ is the optimal strategy in the binary decision problem $\tilde{\mathcal{D}}$, and vice versa.

We turn to explain why we need to assume that $\delta = 1$. This is because for dynamic information acquisition, fixed costs and discounting have different impacts on the decision maker's incentive to run an experiment. To illustrate this, we consider the following example.

u	A	B
a	$1+v$	v
b	v	$1+v$

F	A	B
x_A	p	$1-p$
x_B	$1-p$	p

where $v > 0$ and $p \in (0.5, 1)$. u is the decision maker's utility, F is an experiment, and μ_0 is 0.5.

Suppose the fixed cost $c > 0$ is the only cost of running the experiment F. The decision maker will choose to run F if running F increases her expected utility. If the decision maker runs F and observes x_A, the right action is a. Similarly, if the decision maker observes x_B, the right action is b. After simple calculation, the condition for the decision maker to conduct F is:

$$p - \frac{1}{2} \geq c.$$

As the above inequality indicates, v has no impact on the decision maker's choice concerning whether she runs F or not. When the utility is quasi-linear in fixed costs, the decision maker is interested in the change of her expected utility.

Now, let us think about what happens if the discount factor δ is the only cost of waiting. That is, the decision maker does not incur c when running F. In this case, the condition under which the decision maker chooses F is:

$$\delta\,(p+v) \geq \frac{1}{2} + v.$$

Equivalently,

$$\delta p - \frac{1}{2} \geq (1-\delta)v.$$

Clearly, the decision maker is not willing to conduct the experiment if v is very large. In fact, given δ, there exists v such that the decision maker chooses not to conduct the experiment F. If the decision maker discounts the total amount of utility, there are cases in which she is more interested in the magnitude of expected utility, rather than the changes in her expected utility. For instance, suppose that the payoff matrix u represents monetary values and that v is one billion dollars. Then, the decision maker would choose to take the right action immediately even if she can know the true state in the next period. The decision maker does not choose to run the experiment to receive one more dollar.

The above example shows that if discounting is the cost of experimentation, a constant number cannot be added to or subtracted from a payoff matrix without affecting the decision maker's incentive to conduct an experiment. This leaves us with a trivial normalization of payoffs when discounting is the only cost of experimentation. We can always divide payoffs by some positive number. This normalization barely simplifies analyses. So, if the discounting is the only cost of experimentation, we will consider the class of binary decision problem. However, we will rule out pathological cases. For instance, suppose that all entries of a payoff matrix are negative. In this case, the decision maker is certainly willing to run an experiment even if the experiment is completely uninformative. The motivation of experimentation is not to make an informed decision, but to receive negative utilities later. This is not an interesting case, and we do not analyze such cases.

To sum it up, if fixed cost is the only cost of experimentation, we will rely on Lemma 4.3.1, and analyses will be focused on a subclass of binary decision problems. If discounting is the only cost of experiment, Lemma 4.3.1 cannot be applied, and we will consider the whole class of binary decision problems except uninteresting cases. Clearly, if the decision maker incurs the fixed cost of experimentation and discount future payoffs, we cannot use Lemma

4.3.1.

4.4 Fixed Cost of Experimentation: First Class of Examples

In this section, we assume that $\delta = 1$. That is, we restrict our attention to $\mathscr{D}_{(u,c)} = \{(u, c, 1)\}$, the set of binary decision problems with $\delta = 1$. We consider the set of experiments below:

F	A	B
x_A	p	$1-p$
x_B	$1-p$	p

G	A	B
y_A	q	$1-q$
y_B	$1-q$	q

where $0.5 < q < p < 1$.

Since $\delta = 1$, Lemma 4.3.1 allows us to safely focus on the following payoff matrices:

u_β	A	B
a	1	0
b	0	β

where $\beta > 0$.

For the class of experiments described above and the set of all binary decision problems, we analyze optimal strategies. The first step is to understand the criterion for choosing the right action given a belief. When the decision maker has a belief μ and has to take an action, her right action is a if and only if[13]

$$\mu \geq (1-\mu)\beta.$$

In other words, the decision maker chooses a if and only if her posterior belief is greater than or equal to the threshold belief $\mu^* = \frac{\beta}{1+\beta}$.

[13]We assume that the decision maker chooses a if she is indifferent between a and b.

Let $\mu(\cdot)$ the posterior belief. For the class of experiments that we consider in this section, the posterior belief is ordered in the following way:

$$\mu(x_B, y_B) < \mu(x_B, y_A) < \mu(x_A, y_B) < \mu(x_A, y_A).$$

Also, the prior μ_0 is between $\mu(x_B, y_A)$ and $\mu(x_A, y_B)$. Note that the magnitude of the posterior belief is implied by the prior μ_0 and that μ^* is determined by β. So, the relationship between μ_0 and β can be replaced by the relationship between the posterior beliefs and μ^*. We are going to analyze optimal strategies using μ^* and the posterior beliefs.

Proposition 4.4.1. *In the following cases, the optimal strategy for the decision maker is to take the right action immediately for all $c > 0$.*

1. *If $\mu^* \leq \mu(x_B, y_B)$, the right action is a.*

2. *If $\mu(x_A, y_A) \leq \mu^*$, the right action is b.*

Proposition 4.4.1 describes cases in which the decision maker has an extreme prior. In this case, her belief will not change enough to change the optimal action, regardless of experiment outcomes.

Proposition 4.4.2. *Suppose $\mu(x_B, y_A) \leq \mu^* \leq \mu(x_A, y_B)$.*

1. *Suppose $\mu_0 \leq \mu^* \leq \mu(x_A, y_B)$.*

   - *If $c \leq \mu_0 p - (1 - \mu_0)\beta(1 - p)$, the optimal strategy is to run F. Choose a if x_A is observed and b if x_B is observed.*

   - *If $c > \mu_0 p - (1 - \mu_0)\beta(1 - p)$, the optimal strategy is to take action b.*

2. *Suppose $\mu(x_B, y_A) \leq \mu^* < \mu_0$.*

- *If $c \leq -\mu_0(1-p) + (1-\mu_0)\beta p$, the optimal strategy is to run F. Choose a if x_A is observed and b if x_B is observed.*
- *If $c > -\mu_0(1-p) + (1-\mu_0)\beta p$, the optimal strategy is to take action a.*

Proposition 4.4.2 applies when the outcome of F pins down the right action. If μ^* is between $\mu(x_B, y_A)$ and $\mu(x_A, y_B)$, the outcome of G has no impact on the decision maker's action. For instance, suppose that the decision maker observed x_A. Even if she runs G, she will choose a regardless of the outcome of G. Put differently, if she runs G first and F later, she will ignore the realization of G. Therefore, the decision maker has no incentive to run G. That is, the optimal strategy is to run F if and only if the marginal expected payoffs are greater than the fixed costs.

Two cases are remaining. For those cases, we partially analyze optimal strategies. Let us consider the case in which $\mu(x_A, y_B) < \mu^* < \mu(x_A, y_A)$.

Proposition 4.4.3. *Suppose $\mu(x_A, y_B) < \mu^* < \mu(x_A, y_A)$.*

1. *Given $\mu_0 > \frac{1}{2}$, there exists $\overline{c}$ such that for all $c \in (0, \overline{c})$, the optimal strategy is:*
    - *to run G first. If y_A is observed, run F. If y_B is observed, choose b.*
    - *If (y_A, x_A) is observed, choose a. If (y_A, x_B) is observed, choose b.*
2. *Given $\mu_0 \leq \frac{1}{2}$, there exists $\overline{c}$ such that for all $c \in (0, \overline{c})$, the optimal strategy is:*
    - *to run F first. If x_A is observed, run G. If x_B is observed, choose b.*
    - *If (x_A, y_A) is observed, choose a. If (x_A, y_B) is observed, choose b.*

The reason that Proposition 4.4.3 is a partial characterization is because we do not explicitly compute $\overline{c}$. The exact computation of $\overline{c}$ would be interesting, but we would like to focus on cases in which the decision maker runs the two experiments.

The assumption that $\mu(x_A, y_B) < \mu^* < \mu(x_A, y_A)$ means that the impact of β on the right action dominates the impact of μ_0 on the right action. The assumption means that the decision maker's current right action is b, and she changes her action only if she observes x_A and y_A. Given sufficiently small costs c, the decision maker is willing to run a second experiment, regardless of whether she run F first or G first. In this case, the decision maker will choose a with the probability of observing x_A and y_A and choose b with the complementary probability. Clearly, if we forget costs, the expected payoffs are the same, regardless of the order in which to run experiments. This in turn means that maximizing expected payoffs are equivalent to minimizing costs. Since the expected costs are proportional to the probability of running a second experiment, the decision maker wants to run G first if the probability of observing y_A is lower than the probability of observing x_A. When $\mu_0 > \frac{1}{2}$, the decision maker can save costs by running G first because $q < p$. However, if $\mu_0 < \frac{1}{2}$, the decision maker needs to run F first to reduce the expected costs.

The analysis for the last case is similar to Proposition 4.4.3.

Proposition 4.4.4. *Suppose* $\mu(x_B, y_B) < \mu^* < \mu(x_B, y_A)$.

1. *Given* $\mu_0 < \frac{1}{2}$, *there exists* $\overline{c}$ *such that for all* $c \in (0, \overline{c})$, *the optimal strategy is:*

    - *to run G first. If y_B is observed, run F. If y_A is observed, choose a.*
    - *If (y_B, x_A) is observed, choose a. If (y_B, x_B) is observed, choose b.*

2. *Given $\mu_0 \geq \frac{1}{2}$, there exists $\overline{c}$ such that for all $c \in (0, \overline{c})$, the optimal strategy is:*

   - *to run F first. If x_B is observed, run F. If x_A is observed, choose a.*
   - *If (x_B, y_A) is observed, choose a. If (x_B, y_B) is observed, choose b.*

Proof. Omitted. $\square$

Propositions 4.4.3 and 4.4.4 mean that the decision maker does not always prefer FG to GF. Depending on the binary decision problem, the decision maker is willing to conduct G first in order to save costs. Our analysis fails if $p = q$. So, the following corollary is immediate.

Corollary 4.4.1. *FG is more informative than GF with respect to $\mathscr{D}_{(u,c)}$ if and only if $p = q$.*

We believe that this phenomena can be seen in real life. One example would be a doctor and a patient. Suppose the patient visits the hospital to know whether he has a brain tumor. The doctor has run several tests, and believes that the patient is less likely to have a brain tumor. However, the doctor is a bit uncertain and wants to run two more tests, an MRI scan and an X-ray. Suppose that the doctor is convinced that the patient has a brain tumor only if the MRI scan and the X-ray picture indicate that the patient has a brain tumor. Assuming that the MRI scan is more informative than the X-ray, the doctor takes an X-ray picture of the patient because the doctor believes that the patient is healthy and there is a higher chance that the X-ray picture of the patient shows no brain tumor.

Before we close this section, we would like to mention that the analysis in this section can be applied to general binary experiments. In fact, F and

G do not have to have symmetric information structure. Furthermore, neither of F nor G needs to be more informative than the other. As readers may notice, the argument for optimal strategies does not rely on symmetric information structure and informativeness. In this section, we use the symmetric information for two reasons. First, computation is relatively easy and presentation is compact. Second, we would like to emphasize the fact that one experiment being more informative than the other experiment in static decision problems is not sufficient to guarantee that the more informative experiment is run first in dynamic decision problems.

4.5 Fixed Cost of Experiments: Comparison of FG and GF

In the previous section, we focused on the optimal order in which the decision maker runs experiments, and obtained the negative result, Corollary 4.4.1, that there is no non-trivial comparison of sequential experiments. In this section, we continue to restrict our attention to $\mathscr{D}_{(u,c)}$ and focus on the factors that drive the negative result.

Unlike the previous section, an experiment can have more than two realizations. For $F = (X, f_A, f_B)$, $X = \{x_1, \cdots, x_m\}$, and for $G = (Y, g_A, g_B)$, $Y = \{y_1, \cdots, y_l\}$. We assume that the x_i's and y_j's are ordered in the following way:

$$\frac{f_A(x_1)}{f_B(x_1)} < \cdots < \frac{f_A(x_m)}{f_B(x_m)}$$

and

$$\frac{g_A(y_1)}{g_B(y_1)} < \cdots < \frac{g_A(y_l)}{g_B(y_l)}.$$

This ordering is always possible by relabeling realizations. The ordering means that the posterior belief decreases in the index. For instance, $\mu(x_i)$ is decreasing in i.

To state our result, *decisive experiments* are a useful notion. Roughly speaking, a decisive experiment determines the decision maker's right action with positive probability in some cases. For some prior μ_0 and β, if F is decisive, the decision maker does not need to run G, regardless of how small c is. This means that F being decisive is a necessary condition for FG to be more informative than GF.

As in Section 4, $\mu^*(\beta)$ is the threshold belief, which is given by:

$$\mu^*(\beta) = \frac{\beta}{1+\beta}.$$

Recall that if the belief is larger than or equal to $\mu^*(\beta)$, then action a is optimal. If the belief is smaller than or equal to $\mu^*(\beta)$, then action b is optimal. We define a decisive experiment.

Definition 4.5.1. *F is decisive for the binary decision problem $\mathcal{D}$ and the prior μ_0 if there exist x_a and x_b in X such that*

$$\mu(x_a, y) \geq \mu^*(\beta) \qquad \forall y \in Y,$$

and

$$\mu(x_b, y) < \mu^*(\beta) \qquad \forall y \in Y.$$

In words, if the decision maker observes x_a, then she does not have an incentive to run G since no outcome of G will change her action. Similarly, if the decision maker receives x_b, her right action should be b, regardless of outcomes of G.

It turns out that if there are only two experiments, one experiment is decisive for some binary decision problem, whereas the other experiment cannot be decisive for any binary decision problem.

Lemma 4.5.1. *Suppose* $\frac{f_A(x_1)}{f_B(x_1)}\frac{g_A(y_l)}{g_B(y_l)} > \frac{f_A(x_m)}{f_B(x_m)}\frac{g_A(y_1)}{g_B(y_1)}$*. Then, F is decisive for some binary decision problem $\mathcal{D}$ and some prior μ_0, whereas G cannot be decisive for any binary decision problem and any prior.*

Let us discuss Lemma 4.5.1 and forget the possibility that $\frac{f_A(x_1)}{f_B(x_1)}\frac{g_A(y_l)}{g_B(y_l)} = \frac{f_A(x_m)}{f_B(x_m)}\frac{g_A(y_1)}{g_B(y_1)}$ for a while. If the inequality in Lemma 4.5.1 is flipped, we can relabel two experiments so that the inequality is satisfied. Suppose the inequality is met. Then, we can imagine a circumstance in which for some binary decision problem $\mathcal{D}$ and some prior μ_0,

$$\begin{aligned} \mu(x_1, y_1) \geq \quad \dots \quad &\geq \mu(x_1, y_l) \geq \dots \\ \dots \geq \quad \mu^*(\beta) \quad &\geq \dots \\ \dots \; \mu(x_m, y_1) \geq \quad \dots \quad &\geq \mu(x_m, y_l). \end{aligned}$$

This expression means that a is the right action for the decision maker when she observes x_1. Also, upon receiving x_1, the decision maker has no incentive to run G. Similarly, b is the right action for the decision maker when she observes x_m. This argument show that F is decisive for this binary decision problem.

Let us think about why G cannot be decisive. If G is decisive for some binary decision problem $\mathcal{D}$ and some prior μ_0, then it should be true that at least y_1 and y_l determine the right actions. That is, when the decision maker observes y_1 or y_l, she should have no incentive to run F, no matter how small c is. However, this is impossible. If G is decisive for $\mathcal{D}$ and μ_0, then it must be true that $\mu^*(\beta) \in (\mu(x_m, y_l), \mu(x_1, y_1))$. Otherwise, the decision maker's right action is either a or b, regardless of realizations. If $\mu^*(\beta) \in (\mu(x_m, y_1), \mu(x_1, y_1)$, y_1 alone does not pin down the right action. Clearly, if the decision maker runs F and observes x_1, the right action is a, and if she observes x_m, the right action is b. If $\mu^*(\beta) \in (\mu(x_m, y_l), \mu(x_m, y_1)$,

y_l alone does not determine the right action. If the decision maker runs F and observes x_1, the right action is a, whereas if she observes x_m, the right action is y_l. Therefore, if F is decisive for some binary decision problem and some prior, there is no binary decision problem and no prior for which G is decisive.

One more note about a decisive experiment is that the decisive experiment does not have to be more informative than the other experiment. Look at the following experiments.

F	A	B
x_A	p	$1-q$
x_B	$1-p$	q

G	A	B
y_A	r	$1-r$
y_B	$1-r$	r

where $0.5 < q < p < 1$ and $0.5 < r < 1$. Suppose $\left(\frac{r}{1-r}\right)^2 < \frac{p}{1-p}\frac{q}{1-q}$. One can show that F is decisive for some binary decision problem and that F is not more informative than G in static decision problems.

Let us turn to the possibility that

$$\frac{f_A(x_1)}{f_B(x_1)}\frac{g_A(y_l)}{g_B(y_l)} = \frac{f_A(x_m)}{f_B(x_m)}\frac{g_A(y_1)}{g_B(y_1)}.$$

However, the set of experiments where the inequality in Lemma 4.5.1 becomes an equality has Lebesgue measure zero. Therefore, we may ignore the possibility that the condition in Lemma 4.5.1 is an equality. In other words, we may safely assume that F is decisive for some binary decision problem. This means that when we compare FG and GF, the relationship should be such that $\mathcal{W}_{GF}(\mathcal{D}) \subseteq \mathcal{W}_{FG}(\mathcal{D})$ with respect to $\mathscr{D}$ for all $\mathcal{D} \in \mathscr{D}$. This is obvious because if F is decisive for some binary decision problem $\mathcal{D}$, the decision maker does not have to run G for $\mathcal{D}$.

Using the notion of decisive experiments, we state the result on comparison of FG and GF.

Proposition 4.5.1. *Suppose that F is decisive for some binary decision problem and that $X = \{x_1, x_2\}$. Then, FG is more informative than GF with respect to $\mathscr{D}_{(u,c)}$ if and only if $F = G$.*

Proposition 4.5.1 is compatible with Corollary 4.4.1. Proposition 4.4.2 implies that F is decisive for some binary decision problem. In Section 4, F has two realizations. Therefore, Proposition 4.5.1 means that FG is better than GF if and only if F and G are identical.

The idea for the proof of Proposition 4.5.1 is the same as the characterization of optimal strategies in Section 4. For some parameters, the decision maker's right action is a if she does not run any experiment, and she changes her action only if she observes x_2 and y_l. If c is very small, the decision maker will run a second experiment whenever the second experiment increases the expected utility. This implies that the decision maker wants to save costs. Interestingly, this circumstance also occurs when the prior belief is extreme. This means that there are cases in which μ_A is close to either 0 or 1 and the decision maker wants to minimize the expected costs. Since minimizing costs is to choose the lower between the probability of x_2 and the probability of y_l, in order for FG to be better than GF, the probability of y_l is lower than or equal to the probability of x_2 when the prior is close to zero or one. That is, $f_A(x_2) \leq g_A(y_l)$ and $f_B(x_2) \leq g_B(y_l)$.

For some other parameters, the decision maker's right action is b with no experiment, and she changes her action only if she observes x_1 and y_1. By the symmetry, we find that $f_A(x_1) \leq g_A(y_1)$ and $f_B(x_1) \leq g_B(y_1)$. However, those inequalities cannot be strict inequalities because otherwise $1 = f_A(x_1) + f_A(x_2) < g_A(y_1) + g_A(y_l) \leq 1$. Therefore, if F is decisive and has two realizations, and FG is better than GF, F and G must be the identical experiment.

4.6 Fixed Cost of Experimentation: Second Class of Examples

In this section, we continue to focus on $\mathscr{D}_{(u,c)}$ and look for possibilities of positive results. There are two assumptions in Proposition 4.5.1. Since a decisive experiment generically exists due to Lemma 4.5.1, the only way of obtaining a positive result is to relax the assumption that a decisive experiment has two realizations.

We consider a class of the following experiments.

F	A	B
x_1	$(1-r)p$	$(1-r)(1-p)$
x_2	r	r
x_3	$(1-r)(1-p)$	$(1-r)p$

G	A	B
y_1	q	$1-q$
y_2	$1-q$	q

where $0.5 < q < p < 1$ and $r > 0$.

For a subclass of the above experiments, we can obtain a positive result. That is, for some parameters, FG is better than GF. Let us state the result first, and then explain.

Proposition 4.6.1. *Suppose that* $\frac{p}{1-p} > \left(\frac{q}{1-q}\right)^2$ *and* $r \in \left(\frac{p-q}{p}, \max\left\{1-q, \frac{p-q}{p-0.5}\right\}\right)$. *Then, FG is more informative than GF with respect to $\mathscr{D}_{(u,c)}$.*

Proof. In the Appendix. □

We discuss the conditions in Proposition 4.6.1. The first condition that $\frac{p}{1-p} > \left(\frac{q}{1-q}\right)^2$ has two implications. The first implication is that the posterior beliefs are ordered in a way that F is decisive for some binary decision problems. The second implication is that F is not a combination of G and another independent experiment H. The second implication is important because it implies the violation of Greenshtein Theorem 3.2.

We first show that F is decisive in some binary decision problems. See Lemma below.

Lemma 4.6.1. *Suppose* $\frac{p}{1-p} > \left(\frac{q}{1-q}\right)^2$. *Then, the posterior beliefs are ordered in the following way:*

$$\mu(x_3, y_2) < \mu(x_3, y_1) < \mu(x_2, y_2) < \mu(x_2, y_1) < \mu(x_1, y_2) < \mu(x_1, y_1).$$

Lemma 4.6.1 clearly means that F is decisive for some binary decision problems. For instance, if μ^* is between $\mu_(x_2, y_1)$ and $\mu_(x_1, y_2)$, the outcome of F pins down the right action. Therefore, we state this as corollary without the proof.

Corollary 4.6.1. *Suppose* $\frac{p}{1-p} > \left(\frac{q}{1-q}\right)^2$. *Then, F is decisive for some binary decision problems.*

We prove that F is not a combination of G and another experiment H. For given two experiments $G = (Y, g_A, g_B)$ and $H = (Z, h_A, h_B)$, the combination of G and H, $G \times H$ is defined to be $(Y \times Z, g_A h_A, g_B h_B)$.

Proposition 4.6.2. *Suppose* $\frac{p}{1-p} > \left(\frac{q}{1-q}\right)^2$. *There is no experiment* $H = (Z, h_A, h_B)$ *such that F is a combination of G and H.*

Proposition 4.6.2 is important because if F is a combination of G and another independent experiment, then Greenshtein Theorem 3.2 applies and our analysis will be meaningless. Proposition 4.6.2 has a positive implication that if we restrict our attention to a class of decision problems, the comparison of sequential experiments can be made for a larger class of experiments.

We now turn to the second condition in Proposition 4.6.1 that $r > \frac{p-q}{p}$. This means that in order for FG to be better than GF, r should not be close to zero. If r is very small, then F looks like a binary experiment, and we will obtain the same result as in Section 4.

The third condition we look at is that $r < \frac{p-q}{p-0.5}$. This condition means that F is more informative than G in static decision problems. In fact, we need this condition in case the decision maker runs at most one experiment because of large costs. In those cases, if F is more informative than G, the decision maker chooses F. The Lemma below characterizes a necessary and sufficient condition under which F is more informative than G.

Lemma 4.6.2. *F is more informative than G if and only if $r \in \left[0, \frac{p-q}{p-\frac{1}{2}}\right]$.*

The remaining condition is that $r < \max\left\{1-q, \frac{p-q}{p-0.5}\right\}$. Intuitively, r cannot be very large because it is uninformative. In case p is large and q is close to 0.5, r can be close to one and F is still more informative than G. However, if r is very large, then the decision maker receives the uninformative signal realization with high probability. This means that if the decision maker runs F first, there would be a higher chance of running G. So, r should not be very large. In other words, r needs to be bounded from above in order to prevent the decision maker from running a second experiment with large probability.

The intuition for the proof of Proposition 4.6.1 is the following. For the first class of examples, a more informative experiment F frequently induces the decision maker to run the other experiment G. In order to alleviate this, a more informative experiment F now has three realizations. So, because the decision maker receive extreme outcomes x_1 and x_3 with lower probabilities, the more informative experiment F does not frequently trigger the second experimentation. This is good news. However, the new realization x_2 is uninformative, and clearly, the uninformative realization also induces the decision maker to run the other experiment. The uninformative realization should not occur with high probability, otherwise it also often triggers the second experimentation. Therefore, if the decision maker receives the uninformative signal with some low probability, then she would prefer FG to GF in every binary decision problem.

4.7 Delay by Experimentation: First Class of Examples

In this section, we revisit the first class of examples, assuming that $c = 0$ and $\delta \in [0, 1]$. For experiments, we use the same notation as in the first class of examples.

Since δ is not fixed at 1, we cannot use Lemma 4.3.1 to normalize a payoff matrix. Therefore, we use a generic payoff matrix u:

u	A	B
a	$u_A(a)$	$u_B(a)$
b	$u_A(b)$	$u_B(b)$

Recall that $u_A(a) > u_A(b)$ and $u_B(a) < u_B(b)$. We also impose one more assumption that all entries of the payoff matrix are non-negative. This is because we want to eliminate cases in which the decision maker conducts an experiment in order to take an action later. For instance, suppose all entries of the payoff matrix are negative. Even if an experiment is uninformative, the decision maker is willing to run the experiment because the discounted expected payoff is larger than the current expected payoff.

We define another class of binary decision problems: $\mathscr{D}_{(u,\delta)} = \{(u, 0, \delta)\}$ where all entries of u are non-negative. The analysis of this section is focused on $\mathscr{D}_{(u,\delta)}$.

As in Section 4, we analyze optimal strategies. Since the arguments used in Propositions 4.4.1 and 4.4.2 are independent of whether the decision maker incurs fixed costs or discounts the future payoffs, Propositions 4.4.1

and 4.4.2 apply here, too. Therefore, we move onto cases that correspond to Propositions 4.4.3 and 4.4.4. μ^* is the threshold belief:

$$\mu^* = \frac{u_B(b) - u_B(a)}{u_A(a) - u_A(b)}.$$

Proposition 4.7.1. *Suppose* $\mu(x_A, y_B) < \mu^* < \mu(x_A, y_A)$*.*

1. *Given* $\mu_0 \leq \frac{1}{1+u_A(b)/u_B(b)}$*, there exists* $\underline{\delta}$ *such that for all* $\delta \in [\underline{\delta}, 1)$*, the optimal strategy is:*

   - *to run* F *first. If* x_A *is observed, run* G*. If* y_B *is observed, choose* b*.*
   - *If* (x_A, y_A) *is observed, choose* a*. If* (x_A, y_B) *is observed, choose* b*.*

2. *Given* $\mu_0 > \frac{1}{1+u_A(b)/u_B(b)}$*, there exists* $\underline{\delta}$ *such that for all* $\delta \in [\underline{\delta}, 1)$*, the optimal strategy is:*

   - *to run* G *first. If* y_A *is observed, run* F*. If* y_B *is observed, choose* b*.*
   - *If* (x_A, y_A) *is observed, choose* a*. If* (x_A, y_B) *is observed, choose* b*.*

Proposition 4.7.1 has the same implication as Proposition 4.4.3. Let us consider the second case in Proposition 4.7.1. The assumptions in the second case of Proposition 4.7.1 are satisfied when μ^* is large and μ_0 is large. In this case, the decision maker believes that the true state is almost A but the right action is b. If the decision maker runs F first, then there will be a higher chance that she runs G. However, if the decision maker runs G first, she can take the action b "earlier."

Proposition 4.7.2. *Suppose* $\mu(x_B, y_B) < \mu^* < \mu(x_B, y_A)$*.*

1. *Given* $\mu_0 < \frac{1}{1+u_A(a)/u_A(b)}$*, there exists* $\underline{\delta}$ *such that for all* $\delta \in [\underline{\delta}, 1)$*, the optimal strategy is:*

- *to run G first. If y_B is observed, run F. If y_A is observed, choose a.*
- *If (y_B, x_A) is observed, choose a. If (y_B, x_B) is observed, choose b.*

2. *Given $\mu_0 \geq \frac{1}{1+u_A(a)/u_A(b)}$, there exists $\underline{\delta}$ such that for all $\delta \in [\underline{\delta}, 1)$, the optimal strategy is:*
    - *to run F first. If x_B is observed, run F. If x_A is observed, choose a.*
    - *If (x_B, y_A) is observed, choose a. If (x_B, y_B) is observed, choose b.*

Like Propositions 4.4.3 and 4.4.4, the above two Propositions mean that the decision maker does not always prefer FG to GF. Depending on the binary decision problem, the decision maker is willing to conduct G first in order to take the right action earlier. By the same reason as in Section 4, the following corollary is immediate.

Corollary 4.7.1. *FG is more informative than GF with respect to $\mathscr{D}_{(u,\delta)}$ if and only if $p = q$.*

However, if we further restrict our attention to a subclass of binary decision problems, we obtain a positive result. Let $\mathscr{D}_{(u_\beta,\delta)} = \{(u_\beta, 0, \delta)\}$. The following corollary is immediate from Propositions 4.7.1 and 4.7.2.

Corollary 4.7.2. *FG is more informative than GF with respect to $\mathscr{D}_{(u_\beta,\delta)}$.*

Note that when u_β is used, $u_{\beta,A}(b) = u_{\beta,B}(a) = 0$. In this case, the second part of Proposition 4.7.1 implies that running G first is optimal if $\mu_0 > 1$. However, μ_0 cannot be larger than one. Also, the first part of Proposition 4.7.2 means that if $\mu_0 < 0$, running G first is optimal. Clearly, μ_0 cannot be negative. Therefore, if $u_{\beta,A}(b) = u_{\beta,B}(a) = 0$, the informativeness of static decision problems is carried over in dynamic decision problems.

Finally, the analysis in this section can be applied to general binary experiments. Let us consider the following binary experiments.

F	A	B
x_A	p	$1-q$
x_B	$1-p$	q

G	A	B
y_A	r	$1-s$
y_B	$1-r$	s

Propositions 4.7.1 and 4.7.2 still hold under appropriate revision. Corollary 4.7.1 needs to be changed a bit and says that FG is more informative than GF with respect to $\mathscr{D}_{(u,\delta)}$ if and only if $F = G$. Corollary 4.7.2 also needs to be changed and states that FG is more informative than GF with respect to $\mathscr{D}_{(u_\beta,\delta)}$ if and only if $p \geq p$ and $q \geq s$.

We close this section by discussing the discontinuity in comparison of sequential experiments. For simple argument, let us restrict our attention to a symmetric information structure. Corollary 4.4.1 states that when $\delta = 1$. there is no non-trivial comparison of sequential experiments. However, Corollary 4.7.2 implies that there exists a subclass of binary decision problems for which informativeness of static decision problems implies informativeness of dynamic decision problems, i.e., a more informative experiment in static decision problems should be conducted first. This clearly indicates a discontinuity in comparison of sequential experiments. More precisely, we can imagine two sequences of binary decision problems. The first sequence of binary decision problems is $(u_\beta, c_n, 1)_{n\in\mathbb{N}}$ with $c_n > 0$ converging to zero, and the second sequence of binary decision problems is $(u_\beta, 0, \delta_n)_{n\in\mathbb{N}}$ with $\delta_n \in (0,1)$ converging to one. Clearly, the two sequences of binary decision problems converge to the following binary decision problem $(u_\beta, 0, 1)$. However, for every $(u_\beta, c_n, 1)$, there is no non-trivial comparison of sequential experiments whereas for every $(u_\beta, 0, \delta_n)$, a more informative experiment should be run first provided that doing so is beneficial to the decision maker.

4.8 Delay by Experimentation: Comparison of FG and GF

We investigate whether Proposition 4.5.1 holds when fixed costs are replaced by discounting. That is, we ask ourselves whether Proposition 4.5.1 remains true if we replace $\mathscr{D}_{(u,c)}$ with $\mathscr{D}_{(u,\delta)}$. Not surprisingly, we obtain the same result. We use the same notations as in Section 4.5 to state the following proposition.

Proposition 4.8.1. *Suppose that F is decisive for some binary decision problem and that $X = \{x_1, x_2\}$. Then, FG is more informative than GF with respect to $\mathscr{D}_{(u,\delta)}$ if and only if $F = G$.*

Proposition 4.8.1 is basically the same as Proposition 4.5.1, and states that if we consider all binary decision problems, there is no non-trivial comparison of sequential experiments. Propositions 4.8.1 and 4.5.1 together imply that if we consider all binary decision problems, there is little difference between fixed costs and discounting in terms of comparison of sequential experiments.

4.9 Delay by Experimentation: Second Class of Examples

We continue to restrict ourselves to $\mathscr{D}_{(u,\delta)}$ That is, we replace fixed costs with discounting and look for the counterpart of Proposition 4.6.1. The same notations are used as in Section 4.6.

Lemma 4.6.1 is necessary for F to be decisive in some binary decision problems. Also, Lemma 4.6.2 is necessary to guarantee that F is more informative than G. We need to change the lower bound for r in Proposition 4.6.1 and characterize the range of δ as to whether the decision maker runs an experiment or not.

Proposition 4.9.1. *Suppose that* $\frac{p}{1-p} > \left(\frac{q}{1-q}\right)^2$ *and* $r \in \left(\frac{p-q}{p}, \max\left\{1-q, \frac{p-q}{p-0.5}\right\}\right)$. *Then, FG is more informative than GF with respect to* $\mathscr{D}_{(u,\delta)}$.

The conditions in Proposition 4.6.1 are sufficient conditions for FG to be better than GF when the decision maker incurs fixed costs. It turns out that the conditions in Proposition 4.6.1 are also sufficient conditions for FG to be better than GF when discounting is the cost of experimentation.

Note that Propositions 4.6.1 and 4.9.1 share the same assumptions. So, one could think that if the assumptions are satisfied, for the second class of examples, FG is more informative than GF. However, this is not straightforward. Proposition 4.6.1 may not be true if the discount factor δ is strictly less than one. Proposition 4.9.1 may not be true if the fixed cost c of experimentation is strictly larger than zero. Fortunately, we obtain a positive result.

Since discounting is under consideration, we need to maintain the assumption that all entries of a payoff matrix are non-negative. This will be the only restriction on the set of binary decision problems. Let $\mathscr{D}_+ = \{(u, c, \delta)\}$ where all entries of u are non-negative.

Proposition 4.9.2. *Suppose that* $\frac{p}{1-p} > \left(\frac{q}{1-q}\right)^2$ *and* $r \in \left(\frac{p-q}{p}, \max\left\{1-q, \frac{p-q}{p-0.5}\right\}\right)$. *Then, FG is more informative than GF with respect to* $\mathscr{D}_+$.

The reason why Propositions 4.6.1 and 4.9.1 imply Proposition 4.9.2 is the following. Under the maintained assumptions, the proof of Proposition 4.6.1 implies that the decision maker always minimizes the expected costs by running F first. The proof of Proposition 4.9.1 implies that under the assumption, the decision maker always takes the right action earlier by running F first. Therefore, if the decision maker runs F first, she can not only take the right action earlier but also incur lower expected costs.

There is an important implication of Proposition 4.9.2. By restricting our attention to the class of binary decision problems, we have found a class of experiments for which there exists a non-trivial comparison of sequential experiments in Blackwell sense. In addition, for the second class of experiment, informativeness of static decision problems implies informativeness of dynamic decision problems.

4.10 Conclusion

We summarize our analysis. We consider a class of dynamic decision problems in which the decision maker can run at most two experiments and information acquisition is costly. Then, we investigate the optimal order of experimentation, and look for a partial order of sequential experiments. The first half of our analysis is focused on the case where the decision incurs fixed costs and does not discount the future payoffs. The second half of our analysis is focused on the case where the decision does not incur fixed costs but experimentation delays the decision making process.

The first half of our analysis mainly consists of two parts. For the first part, we consider a class of experiments, and show that even though we restrict our attention to a class of decision problems, there is no non-trivial comparison of sequential experiments. Especially, even if one experiment is more informative than the other experiment, there are decision problems in which it is optimal for the decision maker to conduct the less informative experiment first. The idea for this is mainly driven by cost minimization. One example in which the decision maker is willing to run the less informative experiment first is when the cost of experimentation is small, the decision maker's belief is biased, and she changes her action only if both of the two experiments indicates that she needs to change the action. In such situations, the more informative experiment triggers a second experimentation more often than the less informative experiment does.

The second part deals with another class of experiments. For this class of experiments, we show that for some parameters, informativeness in static decision problems also implies informativeness in dynamic decision problems. That is, it is always optimal for the decision maker to run a more informative experiment first unless experimentation is too costly. The main idea for this positive result is that in this case, a more informative experiment triggers less frequently the second experimentation than a less informative experiment does. Unlike the first part, the decision maker incurs smaller expected costs in dynamic decision problems under consideration.

For the second half of our analysis, we replace fixed costs with discounting. Like the first half of the analysis, there are mainly two parts. We use the same class of experiments as in the first part of the first half of the analysis, and obtain the same result. There is no non-trivial comparison of sequential experiments. However, we find that if we restrict our attention of a subclass of the dynamic decision problems, the decision maker always prefers running a more informative experiment.

In the second part of the second half of the analysis, we use the same class of experiments that we use in the second part of the first half of the analysis. We obtain the same result. For the same parameters as in the second part of the first half of the analysis, the decision makers always prefers running a more informative experiment. This also implies that given the parameters, the decision maker's optimal strategy is to run a more informative experiment even when she incurs fixed costs and experimentation delays the decision making process.

In this paper, we consider two independent experiments and analyze the optimal order in which to run the experiments. One way of extending this paper is to study correlated experiments. If two experiments are independent, then no outcome of one experiment is an indicator of how informative the other experiment is. However, if two experiments are correlated, some outcome of one experiment can be informative about how informative the other experiment is. For instance, as mentioned in Börgers *et al.* (2013), the complementarity of experiments may have some impact on the optimal order of experimentation. Also, the complementarity of experiments may have some implication on the comparison of sequential experiments. Another way of extending this work is to consider more than two experiments. If there are more than two experiments, analysis of the optimal order in which to run experiments would be more difficult, but more interesting.

Appendices

Appendix A. Omitted Proofs in Chapter 2

Proof of Lemma 2.4.1: Choose $\omega \in \mathcal{S}(q)$. The ratio of $\mu_i(s^t)$ to $\mu_j(s^t)$ is given by

$$\begin{aligned}\frac{\mu_i(s^t)}{\mu_j(s^t)} &= \frac{\int \prod_{s\in S} q'(s)^{t\hat{q}_{s,t}(\omega)} F_i(dq')}{\int \prod_{s\in S} q'(s)^{t\hat{q}_{s,t}(\omega)} F_j(dq')} \\ &= \frac{\int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_i(dq')}{\int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_j(dq')}\end{aligned}$$

Note that for all $\omega \in \mathcal{S}(q)$, $\hat{q}_t(\omega)$ converges to q. Suppose μ_i is closer to q than μ_j is. If $\mu_i(\mathcal{S}(q)) > 0$ and $\mu_j(\mathcal{S}(q)) = 0$,

$$\begin{aligned}\frac{\mu_i(s^t)}{\mu_j(s^t)} &= \frac{\mu_i(\mathcal{S}(q))\exp\Big(- tD(\hat{q}_t(\omega), q)\Big) + \int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_i(dq')}{\int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_j(dq')} \\ &= \frac{\mu_i(\mathcal{S}(q)) + \int_{\mathrm{supp}(F_i)\setminus\{q\}} \exp\Big(- tD(\hat{q}_t(\omega), q') + tD(\hat{q}_t(\omega), q)\Big) F_i(dq')}{\int_{\mathrm{supp}(F_j)\setminus\{q\}} \exp\Big(- tD(\hat{q}_t(\omega), q') + tD(\hat{q}_t(\omega), q)\Big) F_j(dq')}.\end{aligned}$$

For every $q' \neq q$, there exists T such that for all $t \geq T$, $D(\hat{q}_t(\omega), q') - D(\hat{q}_t(\omega), q) > 0$. Therefore, as t goes to infinity, denominator converges to zero whereas numerator does not.

Let us consider the second condition in the definition of closeness. Suppose $\mu_i(\mathcal{S}(q)) = 0$ and $\mu_j(\mathcal{S}(q)) = 0$. Let

$$\begin{aligned}\epsilon_i &= \inf_{q'\in\mathrm{supp}(F_i)} D(q, q'), \\ \epsilon_j &= \inf_{q'\in\mathrm{supp}(F_j)} D(q, q').\end{aligned}$$

If the second condition is satisfied, $0 \leq \epsilon_i < \epsilon_j$. In this case, there exists

$b > 0$ such that

$$\begin{aligned}
\lim_{t\to\infty} \frac{\mu_i(s^t)}{\mu_j(s^t)} &= \lim_{t\to\infty} \frac{\int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_i(dq')}{\int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_j(dq')} \\
&\geq \lim_{t\to\infty} b \frac{\int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_i(dq')}{\int e^{-t\epsilon_j} F_j(dq')} \\
&= \lim_{t\to\infty} be^{t\epsilon_j} \int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_i(dq')
\end{aligned}$$

Since $\hat{q}_t(\omega)$ converges to q, $D(\hat{q}_t(\omega), q')$ converges to $D(q, q')$. Choose $\epsilon < \min\{\epsilon_i, \epsilon_j - \epsilon_i\}$. For $\omega \in \mathcal{S}(q)$, define

$$B(\omega) = \left\{ q' \in \Delta S \middle| \epsilon_i - \epsilon \leq \lim_{t\to\infty} D(\hat{q}_t(\omega), q') \leq \epsilon_i + \epsilon \right\}.$$

Then, there exists $b' > 0$ such that

$$\begin{aligned}
\lim_{t\to\infty} \frac{\mu_i(s^t)}{\mu_j(s^t)} &\geq \lim_{t\to\infty} be^{t\epsilon_j} \int \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_i(dq') \\
&\geq \lim_{t\to\infty} be^{t\epsilon_j} \int_{B(\omega)\cap \mathrm{supp}(F_i)} \exp\Big(- tD(\hat{q}_t(\omega), q')\Big) F_i(dq') \\
&\geq \lim_{t\to\infty} bb' e^{t(\epsilon_j - \epsilon_i - \epsilon)} F_i \left(B(\omega) \bigcap \mathrm{supp}(f_1) \right) \\
&= \infty.
\end{aligned}$$

This completes the proof. $\square$

Proof of Proposition 2.4.1: First order condition is

$$\frac{\mu_i(s^t)}{\mu_j(s^t)} \frac{u_i'(c_i^\sigma(s^t))}{u_j'(c_j^\sigma(s^t))} = \lambda.$$

Suppose μ_i is closer to $q \in \Delta S$ than μ_j is. Lemma 2.4.1 implies that $\frac{\mu_i(s^t)}{\mu_i(s^t)}$ diverges ν_q-almost surely. Therefore, $\frac{u_i'(c_i(s^t))}{u_j'(c_j(s^t))}$ converges to zero ν_q-almost surely, which means that $u_j'(c_j^\sigma(s^t))$ goes to infinity ν_q-almost surely, which means that $c_j^\sigma(s^t)$ converges to zero ν_q-almost surely. $\square$

Proof of Lemma 2.5.2: ($\Rightarrow$) Choose an allocation $(c_1, \cdots, c_I)$. Suppose there exists a subgame perfect equilibrium σ such that for all $i \in I$, $c_i^\sigma = c_i$. Suppose that player i deviates at the state history $s^t \in \mathcal{F}_\infty$. Note that the best deviation for player i is to make zero transfer to all other players. Let $\tilde{\sigma}$ be the strategy profile when player i makes zero transfer in the state history s^t.

$$\begin{aligned} u_i(c_i^\sigma(s^t)) + \delta V_i(c_i^\sigma, \mu_i|s^t) &\geq u_i(c_i^{\tilde{\sigma}}(s^t)) + \delta V_i(c_i^{\tilde{\sigma}}, \mu_i|s^t) \\ &\geq u_i(e_i(s^t)) + \delta V_i(e_i, \mu_i|s^t). \end{aligned}$$

The first inequality is because σ is a subgame perfect equilibrium. Clearly, $c_i^{\tilde{\sigma}}(s^t) \geq e_i(s^t)$. Lemma 2.5.1 means that $V_i(c_i^{\tilde{\sigma}}, \mu_i|s^t) \geq V_i(e_i, \mu_i|s^t)$. This completes the proof of the "only if" part.

($\Leftarrow$) Suppose such an allocation exists. We can construct a subgame perfect equilibrium. First, let $\tau_i(s^t)$ be player i's net transfer in s^t.

$$\tau_i(s^t) = e_i(s^t) - c_i(s^t).$$

Let $I_+(s^t)$ be the set of players whose net transfer is strictly positive and $I_-(s^t)$ be the set of players whose net transfer is strictly negative. One can choose a transfer scheme under which for $i \in I_+(s^t)$, $\sum_{j \in I_-} \tau_{ij}(s^t) = \tau_i(s^t)$ and for $j \in I_-(s^t)$, $\sum_{i \in I_+} \tau_{ji}(s^t) = -\tau_j(s^t)$.

Based on such a transfer scheme, we construct the following strategy profile σ. If no player has deviated in the past, all players follow the transfer scheme. If player i has deviated, then all other players make zero transfer to everyone in the future. Then, player i's best response is to make zero transfer to all other players. Since

$$u_i(c_i(s^t)) + \delta V_i(c_i, \mu_i | s^t) \geq u_i(e_i(s^t)) + \delta V_i(e_i, \mu_i | s^t),$$

player i has no incentive to deviate. Therefore, this strategy profile is a subgame perfect equilibrium and for all i, $c_i^\sigma = c_i$. $\square$

Proof of Proposition 2.5.1: Suppose not. Then, we can choose a sequence (σ_n, i_n, s^{t_n}) such that $c_{i_n}^{\sigma_n}(s^{t_n})$ converges to zero.

Choose $0 < \epsilon < \min_{i \in I} \min_{s \in S} e_i(s)$. Then, there exists T such that for every $i \in I$,

$$(1 - \delta^T) u_i(\epsilon) + \delta^{T+1} u_i(e(K)) < \min_{s \in S} u_i(e_i(s)).$$

Let $\alpha = \min_{i \in I} \min_{s \in S} \min_{q \in \text{supp}(F_i)} q(s)$.

Since $c_{i_n}^{\sigma_n}(s^{t_n})$ converges to zero as n increases, there exist player i and n such that

$$\frac{u_i'(c_i^{\sigma_n}(s^{t_n}))}{u_{i_n}'(c_{i_n}^{\sigma_n}(s^{t_n}))} = \frac{\lambda_{i_n} + \Lambda_{i_n}(s^{t_n})}{\lambda_i + \Lambda_i(s^{t_n})} \frac{\mu_{i_n}(s^{t_n})}{\mu_i(s^{t_n})} \leq \alpha^T \min_{i' \in I} \frac{u_{i'}'(e(K) - \epsilon)}{u_{i_n}'(\epsilon)}.$$

This means that player i_n's consumption in s^{t_n} is lower than ϵ. To simplify notations, we suppress n for discussion below and replace i_n with j.

Now, we are going to show that player j's expected utility in s^t is lower than her minimax payoffs. This consists of several claims.

Claim 1. It is impossible that SEC(j) is binding in every state history (s^t, s). Otherwise, player j's expected utility is lower than her minimax payoffs in s^t. Therefore, SEC(j) is not binding in some state s in period $t+1$.

Claim 2. If SEC(j) is not binding in some (s^t, s), then $c_j^\sigma(s^t, s) \le \epsilon$. If SEC($i$) is binding and SEC($j$) is not binding in (s^t, s),

$$\frac{\lambda_j + \Lambda_j(s^t, s)}{\lambda_i + \Lambda_i(s^t, s)} < \frac{\lambda_j + \Lambda_j(s^t)}{\lambda_i + \Lambda_i(s^t)}.$$

If both SEC(i) and SEC(j) are not binding in (s^t, s),

$$\frac{\lambda_j + \Lambda_j(s^t, s)}{\lambda_i + \Lambda_i(s^t, s)} = \frac{\lambda_j + \Lambda_j(s^t)}{\lambda_i + \Lambda_i(s^t)}.$$

The first order condition implies that

$$\begin{aligned}
\frac{u_i'(c_i^\sigma(s^t, s))}{u_j'(c_j^\sigma(s^t, s))} &= \frac{\lambda_j + \Lambda_j(s^t, s)}{\lambda_i + \Lambda_i(s^t, s)} \frac{\mu_j(s^t, s)}{\mu_i(s^t, s)} \\
&\le \frac{\lambda_j + \Lambda_j(s^t)}{\lambda_i + \Lambda_i(s^t)} \frac{\mu_j(s^t, s)}{\mu_i(s^t, s)} \\
&= \alpha^T \frac{\mu_i(s^t)}{\mu_j(s^t)} \frac{\mu_j(s^t, s)}{\mu_i(s^t, s)} \min_{i' \in I} \frac{u_{i'}'(e(K) - \epsilon)}{u_j'(\epsilon)} \\
&\le \alpha^{T-1} \min_{i' \in I} \frac{u_{i'}'(e(K) - \epsilon)}{u_j'(\epsilon)}.
\end{aligned}$$

This shows that $c_j^\sigma(s^t, s)$ is lower than ϵ.

Claim 3. For every $t' = 1, \cdots, T$ and every $\tilde{s}^{t'} \in \mathcal{F}_{t'}$, SEC($j$) is binding in $s^t, \tilde{s}^{t'}$ or $c_j^\sigma(s^t, \tilde{s}^{t'}) \le \epsilon$.

This can be proven by induction. Suppose SEC(j) is not binding in (s^t, s). Then, Claim 2 implies that $c_j^\sigma(s^t, s) \le \epsilon$. Then, due to Claim 1, SEC(j) cannot be binding in period $t+2$ with probability one. That is, SEC(j) is not binding in some $s^t, \tilde{s}^2$. The first order condition implies that

$$\begin{aligned}
\frac{u_i'(c_i^\sigma(s^t,\tilde{s}^2))}{u_j'(c_j^\sigma(s^t,\tilde{s}^2))} &= \frac{\lambda_j + \Lambda_j(s^t,\tilde{s}^2)}{\lambda_i + \Lambda_i(s^t,\tilde{s}^2)} \frac{\mu_j(s^t,\tilde{s}^2)}{\mu_i(s^t,\tilde{s}^2)} \\
&\leq \frac{\lambda_j + \Lambda_j(s^t)}{\lambda_i + \Lambda_i(s^t)} \frac{\mu_j(s^t,\tilde{s}^2)}{\mu_i(s^t,\tilde{s}^2)} \\
&\leq \alpha^{T-2} \min_{i' \in I} \frac{u_{i'}'(e(K)-\epsilon)}{u_j'(\epsilon)}.
\end{aligned}$$

This means that $c_j^\sigma(s^t,\tilde{s}^2) \leq \epsilon$.

Combining these claims, we can prove that SEC(j) is violated in s^t. Consider a path that starts from s^t. Due to Claims 1 – 3, c_j^σ is less than or equal to ϵ until SEC(j) is binding. If SEC(j) is not binding along the path until $t+T$ period, the choice of α and T implies that player j's expected utility in s^t conditional on that path is lower than her minimax payoffs. This implies that arriving at s^t, player j's expected utility conditional on every future path is lower than her minimax payoffs, which means that player j's expected utility in s^t is lower than her minimax payoffs. □

Proof of Lemma 2.7.1: Suppose σ is ambiguity averse Pareto efficient. For $i \in I$, let

$$v_i = \inf_{\mu \in \mathscr{P}^\dagger} V_i(c_i^\sigma, \mu).$$

Let $\overline{c^\sigma}_i$ be a certainty equivalent. That is,

$$\frac{1}{1-\delta} u_i(\overline{c^\sigma}_i) = v_i.$$

Note that $(1-\delta)\delta^{t-1}\mu$ can be thought of as a measure on $(\mathcal{S}, \mathcal{F})$. For any $\mu \in \mathscr{P}^\dagger$ and for each $i \in I$,

$$\begin{aligned}\inf_{\mu' \in \mathscr{P}^\dagger} V_i(c_i^\sigma, \mu') \leq V_i(c_i^\sigma, \mu) \\ &= \frac{1}{1-\delta} \sum_{s^t \in \mathcal{F}_\infty} (1-\delta)\delta^{t-1} u_i(c_i^\sigma(s^t))\mu(s^t) \\ &\leq \frac{1}{1-\delta} u_i \left(\sum_{s^t \in \mathcal{F}_\infty} (1-\delta)\delta^{t-1} \mu(s^t) c_i^\sigma(s^t) \right).\end{aligned}$$

This means that

$$\overline{c^\sigma}_i \leq \sum_{s^t \in \mathcal{F}_\infty} (1-\delta)\delta^{t-1} \mu(s^t) c_i^\sigma(s^t).$$

Hence,

$$\sum_{i \in I} \overline{c^\sigma}_i \leq \sum_{s^t \in \mathcal{F}_\infty} (1-\delta)\delta^{t-1} \mu(s^t) e(s^t)$$

for all $\mu \in \mathscr{P}^\dagger$. The infimum of the right hand side is obtained when $\mu = \nu_{q_1}$ where $q_1 = (1, 0, \cdots, 0) \in \Delta S$. That is, $\sum_{i \in I} \overline{c^\sigma}_i \leq e(1)$. If the inequality is a strict inequality, we can increase every player's utility without violating the feasibility condition. Therefore, $\sum_{i \in I} \overline{c^\sigma}_i = e(1)$. Now, we can construct a stationary strategy profile σ^*. For $i \in I$, there exists $x_i : S \longrightarrow \mathbb{R}_+$ such that $x_i(1) = 0$ and for all $s \in S$,

$$\begin{aligned} c_i^{\sigma^*}(s) &= \overline{c^\sigma}_i + x_i(s), \\ \sum_{i \in I} c_i^{\sigma^*}(s) &= e(s). \end{aligned}$$

By construction, σ^* is ambiguity averse Pareto efficient. Since $c_i^{\sigma^*}(1)$ is the lowest consumption of player i,

$$\inf_{\mu \in \mathscr{P}^\dagger} V_i(c_i^{\sigma^*}, \mu) = \frac{1}{1-\delta} u_i(\overline{c^\sigma}_i) = v_i = \inf_{\mu \in \mathscr{P}^\dagger} V_i(c_i^\sigma, \mu).$$

$\square$

Proof of Proposition 2.8.1: Lemmas 2.7.1 means that for every i,

$$\inf_{\mu \in \mathscr{P}} V_i(c^\sigma, \mu) = \frac{1}{1-\delta} u_i(c_i^\sigma(1))$$

because $c_i^\sigma(1) = \min_{s \in S} c_i^\sigma(s)$. Since for all $i \in I$,

$$\inf_{\mu \in \mathscr{P}} V_i(e_i, \mu) = \frac{1}{1-\delta} \min_{s \in S} e_i(s),$$

there exists $\underline{\delta}$ such that for all $i \in I$ and $s^t \in \mathcal{F}_\infty$,

$$u_i(c_i^\sigma(s^t)) + \frac{\underline{\delta}}{1-\underline{\delta}} u_i(c_i^\sigma(1)) \geq u_i(e_i(s^t)) + \frac{\underline{\delta}}{1-\underline{\delta}} \min_{s \in S} e_i(s).$$

Clearly, for all $\delta > \underline{\delta}$, the above inequality holds. $\square$

Proof of Proposition 2.8.2: Since every player has a rectangular set of priors, the set of priors at the state history s^t is obtained by applying Bayes' rule prior-by-prior. We first show that a full insurance is ambiguity averse Pareto efficient at every state history. Choose a strategy profile σ. Let $CE_i(\sigma, s^t)$ be the certainty equivalent consumption to player i at the state history s^t. At the state history s^t,

$$\begin{aligned}
\frac{1}{1-\delta} u_i(CE_i(\sigma, s^t)) &= \inf_{\mu_i \in \mathscr{P}_i} V_i(c_i^\sigma, \mu_i | s^t) \\
&\leq V_i(c_i^\sigma, \nu_q | s^t) \\
&\leq \frac{1}{1-\delta} u_i\left(\mathbb{E}_{\nu_q}[c_i^\sigma | s^t]\right).
\end{aligned}$$

Summing up, we have

$$\sum_{i\in I} CE_i(\sigma, s^t) \leq \sum_{i\in I} \mathbb{E}_{\nu_q}[c_i^\sigma|s^t] = e.$$

Clearly, if the inequality holds strictly at some state history, a full insurance Pareto dominates σ. Therefore, the inequality must be an equality at every state history, which means that a full insurance is ambiguity averse Pareto efficient at every state history.

When σ is the "no-transfer" strategy profile, the above inequality strictly holds. (This is not true if every player receives constant endowment in every state. However, this case is already a full insurance.) That means that one can find a full insurance σ such that for every i,

$$\begin{aligned}\inf_{\mu_i\in\mathscr{P}_i} V_i(e_i, \mu_i|s^t) &\leq V_i(e_i, \nu_q|s^t)\\ &< \inf_{\mu_i\in\mathscr{P}_i} V_i(c_i^\sigma, \mu_i)\\ &= \frac{1}{1-\delta}u_i(c_i^\sigma).\end{aligned}$$

Given the full insurance σ, there exists $\underline{\delta}$ such that for all $i \in I$ and s^t,

$$u_i(c_i^\sigma) + \frac{\delta}{1-\delta}u_i(c_i^\sigma) \geq u_i(e_i(s^t)) + \delta \inf_{\mu_i\in\mathscr{P}_i} V_i(e_i, \mu_i|s^t).$$

This completes the proof. $\square$

Appendix B. Omitted Proofs in Chapter 3

Proof of Lemma 3.5.1: Define a new function $\psi : \Omega \times T \to \mathbb{R}$ using h. Let

$$\vec{\psi}(x,y) = \vec{h}(x,y) - \vec{h}(0,0).$$

So, $\vec{\psi}$ differs from $\vec{h}$ by a constant vector. Note that $\vec{\psi}$ satisfies incentive compatibility since

$$\begin{aligned}
\vec{\psi}(x,y) \cdot \vec{\mu}(x,y) &= \vec{h}(x,y) \cdot \vec{\mu}(x,y) - \vec{h}(0,0) \cdot \vec{\mu}(x,y) \\
&\geq \vec{h}(x',y') \cdot \vec{\mu}(x,y) - \vec{h}(0,0) \cdot \vec{\mu}(x,y) \\
&= \vec{\psi}(x',y') \cdot \vec{\mu}(x,y).
\end{aligned}$$

For (x,y) and (x',y'), incentive compatibility condition says

$$\begin{aligned}
\vec{\psi}(x,y) \cdot \vec{\mu}(x,y) &\geq \vec{\psi}(x',y') \cdot \vec{\mu}(x,y), \\
\vec{\psi}(x',y') \cdot \vec{\mu}(x',y') &\geq \vec{\psi}(x,y) \cdot \vec{\mu}(x',y'),
\end{aligned}$$

which results in the following inequality.

$$[\vec{\psi}(x,y) - \vec{\psi}(x',y')] \cdot [\vec{\mu}(x,y) - \vec{\mu}(x',y')] \geq 0.$$

When $(x',y') = (0,0)$, we have $\vec{\psi}(x,y) \cdot (\vec{\mu}(x,y) - \vec{\mu}_0) \geq 0$. By the assumption, the inequality implies $\vec{\psi}(x,y) \cdot \hat{\lambda}(y) \geq 0$ for all x, y.

Let us consider an infinitesimal change in x. That is, $x' = x + dx$ and $y' = y$. First order approximation tells us

$$\frac{\partial \vec{\psi}(x,y)}{\partial x} \cdot \frac{\partial \vec{\mu}(x,y)}{\partial x} \geq 0.$$

Now, we are ready to show the claim.

$$\begin{aligned}\frac{d}{dx}\mathbb{E}[\psi|x] &= \int \vec{\psi}(x,y)\cdot\frac{\partial}{\partial x}\vec{\mu}(x,y)dy\\ &= \int \frac{\partial\xi(x,y)}{\partial x}\vec{\psi}(x,y)\cdot\hat{\lambda}(y)dy\\ &\geq 0.\end{aligned}$$

It is straightforward to show that the second order derivative is non-negative.

$$\begin{aligned}\frac{d^2}{dx^2}\mathbb{E}[\psi|x] &= \int \frac{\partial}{\partial x}\vec{\psi}(x,y)\cdot\frac{\partial}{\partial x}\vec{\mu}(x,y)+\vec{\psi}(x,y)\cdot\frac{\partial^2}{\partial x^2}\vec{\mu}(x,y)dy\\ &= \int \frac{\partial}{\partial x}\vec{\psi}(x,y)\cdot\frac{\partial}{\partial x}\vec{\mu}(x,y)+\frac{\partial^2\xi(x,y)}{\partial x^2}\vec{\psi}(x,y)\cdot\hat{\lambda}(y)dy\\ &\geq 0.\end{aligned}$$

The final step is to realize that

$$\begin{aligned}\mathbb{E}[\psi|x] &= \mathbb{E}[h|x]-\vec{h}(0,0)\cdot\mathbb{E}[\vec{\mu}(x,y)|x]\\ &= \mathbb{E}[h|x]-\vec{h}(0,0)\cdot\mu_0.\end{aligned}$$

Since $\mathbb{E}[h|x]$ is $\mathbb{E}[\psi|x]$ plus some constant, it is clear that $\frac{d}{dx}\mathbb{E}[h|x]\geq 0$ and $\frac{d^2}{dx^2}\mathbb{E}[h|x]\geq 0$.

□

Proof of Lemma 3.5.2: Since $\mathbb{E}[h|x]$ is convex in x and $z(x)$ is strictly concave in x, a function $\eta(x)=\mathbb{E}[h|x]-z(x)$ is also convex in x. Since $\mathbb{E}[h|\underline{x}]=z(\underline{x})$ and $\mathbb{E}[h|\overline{x}]=z(\overline{x})$, the number of points where $\mathbb{E}[h|x]$ and $z(x)$ intersect is equal to the number of roots of $\eta(x)$. Since $\eta''>0$, if η crosses the horizontal axis from below, it never crosses horizontal axis again. Since the sign of η' changes at most once, η has at most two roots. □

Proof of Proposition 3.6.1:

1. This is the principal's outside option. If agents' outside options are always greater than the principal's utility, she is not willing to hire an agent.

2. When the principal admits all types, then $\mathbb{E}[h|x]$ should be tangent to $z(x)$ at some point. Define

$$g = \frac{z'(x^*)}{\frac{d}{dx}\mathbb{E}[h|x^*]}(h - \mathbb{E}[h|x^*]) + z(x^*).$$

Note that

$$\begin{aligned}\mathbb{E}[g|x^*] &= z(x^*),\\ \frac{d}{dx}\mathbb{E}[g|x^*] &= z'(x^*).\end{aligned}$$

Therefore, g is tangent to z at point x^*. We want to find an optimal x^*. Let

$$\alpha(x^*) = \frac{z'(x^*)}{\frac{d}{dx}\mathbb{E}[h|x^*]}.$$

Let us calculate the first order condition w.r.t. x^*.

$$\begin{aligned}\frac{d}{dx^*}\mathbb{E}[g] &= \alpha'(x^*)(\mathbb{E}[h] - \mathbb{E}[h|x^*]) - \alpha(x^*)\frac{d}{dx^*}\mathbb{E}[h|x^*] + z'(x^*)\\ &= \alpha'(x^*)(\mathbb{E}[h] - \mathbb{E}[h|x^*]).\end{aligned}$$

Note that $\alpha' < 0$. Therefore, if $\mathbb{E}[h] > \mathbb{E}[h|x^*]$, we can lower the expected payment by choosing a higher x^*. Similarly, if $\mathbb{E}[h] < \mathbb{E}[h|x^*]$, we can lower the expected payment by choosing a lower x^*. Hence, if x^* is already optimal, $\mathbb{E}[h|x^*] = \mathbb{E}[h] = z(x^*)$. This means that $x^* = z^{-1}(\mathbb{E}[h])$.

3. When the principal wants to admit only types higher than $\overline{x}$, $\mathbb{E}[h|x = 0] = z(0)$. If $\mathbb{E}[h|x = 0] < z(0)$, then the principal can find a better contract by rotating h around $\overline{x}$ so that she can pay less to those who accept a contract without changing $\overline{x}$. The equation can be derived from case 5, and we will discuss it at the end of the proof.

4. The argument is symmetric.

5. In this case, types lower than $\underline{x}$ and higher than $\overline{x}$ accept a contract. If this is optimal, the principal should have no incentive to change those thresholds. Expected benefit to the principal is

$$\begin{aligned}\mathbb{E}[u-h] = &\int_0^{\underline{x}} dxdy[\vec{u}(a^*(\vec{\mu}(x,y))) - \vec{h}(x,y)] \cdot \vec{\mu}(x,y)p(x,y) \\ &+ \int_{\overline{x}}^1 dxdy[\vec{u}(a^*(\vec{\mu}(x,y))) - \vec{h}(x,y)] \cdot \vec{\mu}(x,y)p(x,y).\end{aligned}$$

Let

$$\begin{aligned}H_1(h,\underline{x}) &= \int_0^{\underline{x}} dxdy[\vec{u}(a^*(\vec{\mu}(x,y))) - \vec{h}(x,y)] \cdot \vec{\mu}(x,y)p(x,y), \\ H_2(h,\overline{x}) &= \int_{\overline{x}}^1 dxdy[\vec{u}(a^*(\vec{\mu}(x,y))) - \vec{h}(x,y)] \cdot \vec{\mu}(x,y)p(x,y).\end{aligned}$$

Calculus of variation will be used. When h changes to $h+\delta h$, $\underline{x}$ changes to $\underline{x}+\underline{\delta x}$ and $\overline{x}$ changes to $\overline{x}+\overline{\delta x}$. First, we calculate how H_1 changes.

$$\begin{aligned}&H_1(h+\delta h, \underline{x}+\underline{\delta x}) - H_1(h,\underline{x}) \\ &= \int_0^{\underline{x}+\underline{\delta x}} [\vec{u} - \vec{h} - \overrightarrow{\delta h}] \cdot \vec{\mu}p(x)dxdy - \int_0^{\underline{x}} [\vec{u} - \vec{h}] \cdot \vec{\mu}p(x)dxdy \\ &= -\int_0^{\underline{x}+\underline{\delta x}} \overrightarrow{\delta h} \cdot \vec{\mu}p(x)dxdy + \int_{\underline{x}}^{\underline{x}+\underline{\delta x}} [\vec{u} - \vec{h}] \cdot \vec{\mu}p(x)dxdy \\ &\approx -\mathbb{E}[\delta h | x \leq \underline{x}]P(\underline{x}) + \underline{\delta x}\mathbb{E}[u-h|\underline{x}]p(\underline{x}).\end{aligned}$$

Note that

$$\mathbb{E}[h+\delta h|\underline{x}+\underline{\delta x}] = z(\underline{x}+\underline{\delta x}).$$

Up to the first order,

$$\mathbb{E}[h|\underline{x}] + \frac{d}{dx}\mathbb{E}[h|\underline{x}]\underline{\delta x} + \mathbb{E}[\delta h|\underline{x}] = z(\underline{x}) + \frac{dz(\underline{x})}{dx}\underline{\delta x},$$

which results in

$$\underline{\delta x} = -\frac{\mathbb{E}[\delta h|\underline{x}]}{\frac{d}{dx}\mathbb{E}[h|\underline{x}] - \frac{dz(\underline{x})}{dx}}.$$

So, change in H_1 is

$$\begin{aligned}
&H_1(h+\delta h, \underline{x}+\underline{\delta x}) - H_1(h, \underline{x}) \\
&= -\mathbb{E}[\delta h|x \le \underline{x}]F(\underline{x}) - \frac{\mathbb{E}[\delta h|\underline{x}]}{\frac{d}{dx}\mathbb{E}[h|\underline{x}] - \frac{dz(\underline{x})}{dx}}\mathbb{E}[u-h|\underline{x}]p(\underline{x}).
\end{aligned}$$

Change in H_2 is similar.

$$\begin{aligned}
&H_2(h+\delta h, \overline{x}+\overline{\delta x}) - H_2(h, \overline{x}) \\
&= \int_{\overline{x}+\overline{\delta x}}^{1} [\vec{u} - \vec{h} - \overrightarrow{\delta h}] \cdot \vec{\mu} p(x) dx dy - \int_{\overline{x}}^{1} [\vec{u} - \vec{h}] \cdot \vec{\mu} p(x) dx dy \\
&= -\int_{\overline{x}+\overline{\delta x}}^{1} \overrightarrow{\delta h} \cdot \vec{\mu} p(x) dx dy - \int_{\overline{x}}^{\overline{x}+\overline{\delta x}} [\vec{u} - \vec{h}] \cdot \vec{\mu} p(x) dx dy \\
&\approx -\mathbb{E}[\delta h|x \ge \overline{x}]\left(1 - P(\overline{x})\right) - \overline{\delta x}\mathbb{E}[u-h|\overline{x}]p(\overline{x}).
\end{aligned}$$

Calculation of $\overline{\delta x}$ is similar to that of $\underline{\delta x}$, and we have

$$\overline{\delta x} = -\frac{\mathbb{E}[\delta h|\overline{x}]}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{dz(\overline{x})}{dx}}.$$

So, change in H_2 is

$$
\begin{aligned}
&H_2(h+\delta h, \overline{x}+\overline{\delta x}) - H_2(h,\overline{x}) \\
&\approx -\mathbb{E}[\delta h|x \geq \overline{x}]\,(1-P(\overline{x})) + \frac{\mathbb{E}[\delta h|\overline{x}]}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{dz(\overline{x})}{dx}}\mathbb{E}[u-h|\overline{x}]p(\overline{x}).
\end{aligned}
$$

Hence,

$$
\begin{aligned}
&\mathbb{E}[u-(h+\delta h)] - \mathbb{E}[u-h] \\
&= -\mathbb{E}[\delta h|x \leq \underline{x}]P(\underline{x}) - \frac{\mathbb{E}[\delta h|\underline{x}]}{\frac{d}{dx}\mathbb{E}[h|\underline{x}] - \frac{dz(\underline{x})}{dx}}\mathbb{E}[u-h|\underline{x}]p(\underline{x}) \\
&\quad - \mathbb{E}[\delta h|x \geq \overline{x}](1-P(\overline{x})) + \frac{\mathbb{E}[\delta h|\overline{x}]}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{dz(\overline{x})}{dx}}\mathbb{E}[u-h|\overline{x}]p(\overline{x}).
\end{aligned}
$$

Suppose h is optimal. A deviation should not be profitable. First deviation we consider is $(1 \pm \epsilon)h \mp \epsilon z(\underline{x})$. That is, $\delta h = \pm\epsilon(h - z(\underline{x}))$. Since $\mathbb{E}[\delta h|\underline{x}] = 0$,

$$
\begin{aligned}
&\frac{\mathbb{E}[h - z(\underline{x})|\overline{x}]}{\frac{d}{dx}\mathbb{E}[h|\overline{x}] - \frac{dz(\overline{x})}{dx}}\mathbb{E}[u-h|\overline{x}]p(\overline{x}) \\
&= \mathbb{E}[h - z(\underline{x})|x \leq \underline{x}]P(\underline{x}) + \mathbb{E}[h - z(\underline{x})|x \geq \overline{x}](1-P(\overline{x}))
\end{aligned}
$$

Second deviation will be $\pm\epsilon(h - z(\overline{x}))$. This corresponds to $h' = (1 \pm \epsilon)h \mp \epsilon z(\overline{x})$. The second deviation results in

$$
\begin{aligned}
&-\frac{\mathbb{E}[h - z(\overline{x})|\underline{x}]}{\frac{d}{dx}\mathbb{E}[h|\underline{x}] - \frac{dz(\underline{x})}{dx}}\mathbb{E}[u-h|\underline{x}]p(\underline{x}) \\
&= \mathbb{E}[h - z(\overline{x})|x \leq \underline{x}]P(\underline{x}) + \mathbb{E}[h - z(\overline{x})|x \geq \overline{x}](1-P(\overline{x}))
\end{aligned}
$$

Finally, we obtain the result by replacing $\mathbb{E}[h|\underline{x}]$ with $z(\underline{x})$ and $\mathbb{E}[h|\overline{x}]$ with $z(\overline{x})$.

For case 3, we use the first equation to find an optimal $\overline{x}$ with $\underline{x} = 0$. For case 4, we use the second equation and plug $x = 1$ in $\overline{x}$. $\square$

Proof of Lemma 3.7.1: Choose an incentive compatible contract g. Let $\underline{x}$ and $\overline{x}$ be the corresponding two thresholds. Through an affine transformation, we can make $\mathbb{E}[h|x]$ pass through $\underline{x}$ and $\overline{x}$. This contract costs the principal less than g. A similar argument can be applied when $\mathbb{E}[g|x]$ is tangent to $z(x)$. $\square$

Proof of Proposition 3.7.1: Recall that the expected payment is convex in x. So, if expected payment is linear in x, then it is a best contract. Let h be an incentive compatible contract. Since $\xi_X(x)$ can replace x via re-parameterization, we can assume that $\xi_X(x) = x$. Then,

$$\frac{d}{dx}\mathbb{E}[h|x] = \int \xi_Y(y)\vec{h}(x, y) \cdot \hat{\lambda}(y)dy.$$

If $\vec{h}(x, y)$ is independent of x, then the above derivative is constant, and $\mathbb{E}[h|x]$ linearly increases in x. We will construct an incentive compatible contract.

Define

$$\vec{v}(y) = \hat{\lambda}(y) - (\hat{\lambda}(y) \cdot \hat{\mu}_0)\hat{\mu}_0.$$

Choose $\vec{h}(y) = \hat{v}(y)$. We need to show that $\vec{h}(y)$ satisfies the incentive compatibility condition. If an agent with a type (x, y) reports truthfully, he expects to receive the amount below.

$$\begin{aligned}\vec{h}(y)\cdot\vec{\mu}(x,y) &= \hat{v}(y)\cdot(\vec{\mu}_0 + x\xi_Y(y)\hat{\lambda}(y))\\ &= x\xi_Y(y)\hat{v}(y)\cdot\hat{\lambda}(y)\\ &= x\xi_Y(y)\hat{v}(y)\cdot\vec{v}(y)\\ &= x\xi_Y(y)|\vec{v}(y)|.\end{aligned}$$

When he reports (x', y'), expected payment to him is

$$\begin{aligned}\vec{h}(y')\cdot\vec{\mu}(x,y) &= \hat{v}(y')\cdot(\vec{\mu}_0 + x\xi_Y(y)\hat{\lambda}(y))\\ &= x\xi_Y(y)\hat{v}(y')\cdot\hat{\lambda}(y)\\ &= x\xi_Y(y)\hat{v}(y')\cdot\vec{v}(y)\\ &= x\xi_Y(y)|\vec{v}(y)|(\hat{v}(y')\cdot\hat{v}(y)).\end{aligned}$$

Since $\hat{v}(y')\cdot\hat{v}(y)\leq 1$, h is incentive-compatible.

However, $\vec{h}=\hat{v}(y)$ does not rule out the possibility that an agent reports a different agent type x' since h does not depend on x. In order to have have a strong incentive compatibility condition, we can add $\epsilon(\hat{\mu}(x,y)-\hat{\mu}_0)$ to $\vec{h}$. Then,

$$\mathbb{E}[h|x] = x\int \xi_Y(y)|v(y)|dy + \epsilon\mathbb{E}[\hat{\mu}-\hat{\mu}_0|x].$$

As ϵ goes to zero, $\mathbb{E}[h|x]$ converges to a linear function of x. $\square$

Proof of Proposition 3.7.2: We are going to rely on Radner and Stiglitz (1984) result (Theorem 1, p. 36), and we need to convert our environment to Radner and Stiglitz's environment.

Let h be a linear incentive compatible contract, and S be the corresponding proper scoring rule. The agent in our model corresponds to a decision maker in Radner and Stiglitz environment. The set of actions is

$M = \{\vec{\mu}(x,y)|(x,y) \in X \times Y\}$, and the agent chooses a belief in M. $S(\omega, \vec{\mu})$ is the agent's utility when he reports $\vec{\mu} \in M$ and $\omega \in \Omega$ is realized. $x \in X$ corresponds to the parameter that index a family of information structure in Radner and Stiglitz (1984).

Since S is associated with a linear incentive compatible contract, $\frac{d}{dx}\mathbb{E}[S|x] > 0$ for all $x \in X$. This is incompatible with Radner and Stiglitz (1984) result, and therefore, some assumptions made in Radner and Stiglitz (1984) must be violated.

There are five assumptions used in Radner and Stiglitz (1984). We list the assumptions.

1. For every $\omega \in \Omega$ and $x \in X$, $f(\cdot|x, \omega)$ is differentiable at $x = 0$.
2. For every $\omega \in \Omega$ and $\vec{\mu} \in M$, $S(\omega, \vec{\mu})$ is monotone non-increasing in x.
3. For every $\omega \in \Omega$ and $\vec{\mu} \in M$, $g(\vec{\mu}, x)$ is monotone non-increasing in x.
4. For every $\omega \in \Omega$, $S(\omega, \cdot)$ is continuous on $M \times X$.
5. A decision function is flat and continuous at $x = 0$.

The first assumption is satisfied because of Assumption 1. The second assumption is automatically satisfied because S is independent of x. This is because the principal cannot observe the agent type and the signal realization. The function g in the third assumption is a restriction on actions available to the agent. In our model, the agent's report is not restricted by his type. So, the third assumption is satisfied. Let us postpone the fourth assumption and jump to the fifth assumption.

In Radner and Stiglitz (1984), a decision function is said to be flat if it results in the same action regardless of signal realizations. In our model, the agent's action is to report the true posterior belief. That is, when the agent draws the type x and observes the signal realization y, the optimal action is to report $\vec{\mu}(x,y)$. Recall that $\vec{\mu}(0,y)=\vec{\mu}_0$ for all $y\in Y$. Therefore, $\vec{\mu}(\cdot,y)$ is flat at $x=0$. Due to Assumption 1, $\vec{\mu}(\cdot,y)$ is continuous at $x=0$. Therefore, $\vec{\mu}(\cdot,\cdot)$ is flat and continuous at $x=0$, which means that the fifth assumption is satisfied.

The above arguments implies that the fourth assumption has to be violated. That is, for some $\omega\in\Omega$, $S(\omega,\cdot)$ is discontinuous at some $\vec{\mu}\in M$. We need to show that the prior is one of points at which S is discontinuous. Suppose not. Proposition 1 implies that there exist a convex function $\phi: M^{co}\longrightarrow\mathbb{R}$ and a subgradient of ϕ, $\vec{G}: M^{co}\longrightarrow\mathbb{R}^n$ such that

$$\vec{S}(\vec{\mu})=\phi(\vec{\mu})-\vec{G}(\vec{\mu})\cdot\vec{\mu}+\vec{G}(\vec{\mu}).$$

If $\vec{S}$ is continuous at $\vec{\mu}_0$, that means that a gradient of ϕ exists at $\vec{\mu}_0$. In other words, ϕ is differentiable at $\vec{\mu}_0$. So, we can use Taylor expansion to calculate the expected score near $x=0$. Choose a small ε. Up to the first order,

$$\begin{aligned}\mathbb{E}[S|x=\varepsilon] &= \int \phi(\vec{\mu}(\varepsilon,y))dy\\ &\approx \int \phi(\vec{\mu}_0)+\nabla\phi(\vec{\mu}_0)\cdot(\vec{\mu}(\varepsilon,y)-\vec{\mu}_0)\,dy\\ &= 0.\end{aligned}$$

The second term vanishes because the expectation of the posterior beliefs must be the prior. Therefore,

$$\frac{d}{dx}\mathbb{E}[S|x]\Big|_{x=0}.$$

However, this is a contradiction, which completes the proof. □

Appendix C. Omitted Proofs in Chapter 4

Proof of Lemma 4.3.1: ($\Rightarrow$) Suppose σ is optimal in $\mathcal{D}$. Let σ_{FG} be the optimal strategy in Σ_{FG}. Let $\vec{v} = (u_A(b), u_B(a))$. For $\alpha \in A$, $\vec{u}(\alpha) = (u_A(\alpha), u_B(\alpha))$, and $\vec{\mu} = (\mu_A, \mu_B)$. Note that

$$\vec{u}_\beta(a) = \frac{1}{u_A(a) - u_A(b)}(\vec{u}(a) - \vec{v}),$$
$$\vec{u}_\beta(b) = \frac{1}{u_A(a) - u_A(b)}(\vec{u}(b) - \vec{v}).$$

Given the sequence FG, the optimal strategy is simple. If running G is beneficial, the decision maker runs G in period 2. Otherwise, choose the right action. Let us figure out the right action first. Given the binary decision problem $\mathcal{D}$, the decision maker chooses a if and only if

$$\mu_A u_A(a) + (1 - \mu_A) u_B(a) \geq \mu_A u_A(b) + (1 - \mu_A) u_B(b).$$

Rearranging the expression, we have

$$\frac{\mu_A}{1 - \mu_A} \geq \frac{u_B(b) - u_B(a)}{u_A(a) - u_A(b)} = \beta.$$

When the decision maker faces $\tilde{\mathcal{D}}$, she chooses action a if and only if

$$\mu_A \geq (1 - \mu_A)\beta.$$

So, if running G is suboptimal, the decision maker chooses the same action in $\mathcal{D}$ and $\tilde{\mathcal{D}}$.

In $\mathcal{D}$, the decision maker runs G if and only if the expected utility is larger than costs c. Suppose $\vec{\mu}$ is the posterior belief, and let $\alpha^*(\vec{\mu})$ be the optimal action given the posterior belief $\vec{\mu}$.

$$\mathbb{E}[u|\vec{\mu}, \sigma_{FG}] - c \geq \vec{u}(\alpha^*(\vec{\mu})) \cdot \vec{\mu}.$$

Let us investigate whether the same condition holds when the decision maker faces $\tilde{\mathcal{D}}$. That is, we need to show that given $\tilde{\mathcal{D}}$, the decision maker runs G if and only if

$$\mathbb{E}[u_\beta|\vec{\mu}, \sigma_{FG}] - c' \geq \vec{u}_\beta(\alpha^*(\vec{\mu})) \cdot \vec{\mu}.$$

After some calculation, one can show that the left hand side is:

$$\frac{1}{u_A(a) - u_A(b)} \left(\mathbb{E}[u|\vec{\mu}, \sigma_{FG}] - c\right) - \frac{1}{u_A(a) - u_A(b)} \vec{v} \cdot \vec{\mu}.$$

The right hand side is:

$$\frac{1}{u_A(a) - u_A(b)} \vec{u}(\alpha^*(\vec{\mu})) \cdot \vec{\mu} - \frac{1}{u_A(a) - u_A(b)} \vec{v} \cdot \vec{\mu}.$$

Hence,

$$\begin{aligned} &\mathbb{E}[u|\vec{\mu}, \sigma_{FG}] - c \geq \vec{u}(\alpha^*(\vec{\mu})) \cdot \vec{\mu} \\ \iff &\mathbb{E}[u_\beta|\vec{\mu}, \sigma_{FG}] - c \geq \vec{u}_\beta(\alpha^*(\vec{\mu})) \cdot \vec{\mu}. \end{aligned}$$

Now, we need to calculate the ex-ante expected utility. Let $\vec{\mu}(\cdot)$ be the posterior belief. In $\tilde{\mathcal{D}}$,

$$
\begin{aligned}
&\mathbb{E}[u_\beta|\sigma_{FG}] \\
&= \sum_{x\in X, \alpha\in A} \vec{u}_\beta(\alpha)\cdot\vec{\mu}(x)\sigma^1_{FG}(\alpha,x)f(x) \\
&\quad + \sum_{x\in X}\sigma^1_{FG}(G,x)f(x)\left(-c+\sum_{y\in Y,\alpha\in A}\vec{u}_\beta(\alpha)\cdot\vec{\mu}(x,y)\sigma^2_{FG}(\alpha,x,y)g(y)\right).
\end{aligned}
$$

Plug $\vec{u}_\beta(\alpha) = \frac{1}{u_A(a)-u_A(b)}(\vec{u}(\alpha)-\vec{v})$ into the above equation. Then,

$$
\mathbb{E}[u_\beta|\sigma_{FG}] = \frac{1}{u_A(a)-u_A(b)}\mathbb{E}[u|\sigma_{FG}] - \frac{1}{u_A(a)-u_A(b)}\vec{v}\cdot\vec{\mu}_0.
$$

Let σ_{GF} be the optimal strategy in Σ_{GF}. The above calculation shows that

$$
\mathbb{E}[u|\sigma_{FG}] \gtrless \mathbb{E}[u|\sigma_{GF}] \Longleftrightarrow \mathbb{E}[u_\beta|\sigma_{FG}] \gtrless \mathbb{E}[u_\beta|\sigma_{GF}]
$$

($\Leftarrow$) The above argument also proves this direction, too. Therefore, if σ is optimal in $\mathcal{D}$ if and only if it is optimal in $\tilde{\mathcal{D}}$. □

Proof of Proposition 4.4.1: Let us explain the first case. In this case, the lowest posterior belief is higher than or equal to the threshold belief. This means that even if the decision maker had run the two experiments, she would have chosen action a. However, the decision maker can always choose a without running an experiment. Since the decision maker does not want to incur costs, she should not run an experiment and take the right action. The same argument can be made for the second case. □

Proof of Proposition 4.4.2: In this case, the decision maker has no incentive to run G at all. This is because if the decision maker observes either x_A or x_B, no outcome of G will change the decision maker's right action. So, what matters is whether the decision maker runs F or not.

If $\mu_0 \leq \mu^* \leq \mu(x_A, y_B)$, the current right action is b. The decision maker wants to run F only if the expected payoffs increase. That is, running F is optimal only if

$$\mu_0 p + (1 - \mu_0) p\beta - c \geq (1 - \mu_0)\beta.$$

This is the proof of the first part. The argument for the second part is similar. □

Proof of Proposition 4.4.3: If $c = 0$, the decision maker will run the two experiments. Since the expected payoffs are continuous in c, there exists $\bar{c}$ such that as long as costs are smaller than $\bar{c}$, the decision maker is willing to run the two experiments.

Let us calculate utility vectors. Suppose the decision maker runs F first, and let σ_{FG} be the optimal strategy.

$$\begin{aligned} v_A(\sigma_{FG}) &= p(-c + q), \\ v_B(\sigma_{FG}) &= \beta p + (1 - p)(-c + \beta q). \end{aligned}$$

Let σ_{GF} be the optimal strategy when the decision maker runs G first.

$$\begin{aligned} v_A(\sigma_{GF}) &= q(-c + p), \\ v_B(\sigma_{GF}) &= \beta q + (1 - q)(-c + \beta p). \end{aligned}$$

Let Δv be the difference between the above two vectors. That is, $\Delta v = v(\sigma_{FG}) - v(\sigma_{GF})$.

$$\Delta v = c(p - q)(-1, 1).$$

If $\mu_0 > 0.5$, $\Delta v \cdot (\mu_0, 1 - \mu_0)$ is weakly negative. This means that running G first is optimal. If $\mu_0 \leq 0.5$, $\Delta v \cdot (\mu_0, 1 - \mu_0)$ is positive, and it is optimal for the decision maker to run F first. $\square$

Proof of Lemma 4.5.1: We can order signal realizations in the following way. For F,

$$\frac{f_A(x_1)}{f_B(x_1)} \geq \cdots \geq \frac{f_A(x_m)}{f_B(x_m)}$$

and for G,

$$\frac{g_A(y_1)}{g_B(y_1)} \geq \cdots \geq \frac{g_A(y_l)}{g_B(y_l)}.$$

The posterior belief given x, y can be written as

$$\mu(A|x,y) = \frac{1}{1 + \frac{1-\mu_A}{\mu_A}\frac{f_B(x)}{f_A(x)}\frac{g_B(y)}{g_A(y)}}.$$

Let us define two sets. For μ_A and $x \in X$,

$$Y_a(\mu_A, x) = \{y \in Y | \mu(A|x,y) \geq \mu^*(\beta)\}.$$

For μ_A and $y \in Y$,

$$X_a(\mu_A, y) = \{x \in X | \mu(A|x,y) \geq \mu^*(\beta)\}.$$

Since $\mu(A|\cdot)$ is increasing in μ_A, for $\mu_A < \mu'_A$, $Y_a(\mu_A, x) \subset Y_a(\mu'_A, x)$ and $X_a(\mu_A, y) \subset X_a(\mu'_A, y)$. Also, if $i > i'$, $Y_a(\mu_A, x_i) \subset Y_a(\mu_A, x_{i'})$ because $\mu(A|x_i, y) \leq \mu(A|x_{i'}, y)$. Similarly, if $j > j'$, $X_a(\mu_A, y_j) \subset X_a(\mu_A, y_{j'})$.

Let ν_{x_1} be the smallest belief such that $Y_a(\nu_{x_1}, x_1) = Y$. Similarly, let ν_{y_1} be the smallest belief such that $X_a(\nu_{y_1}, y_1) = X$. Generically, we can assume that $\nu_{x_1} < \nu_{y_1}$. If they are the same, we can change the information structure so that the strict inequality holds while the new information structure remains arbitrarily close to the original structure. This means that when $\mu_A = \nu_{x_1}$, $\mu(A|x_1, y) \geq \mu^*(\beta)$ for all $y \in Y$. Since $X_a(\nu_{x_1}, y_1) \neq X$, it must be true that $x_m \notin X_a(\nu_{x_1}, y_1)$, which means that $\mu(A|x_m, y_1) < \mu^*(\beta)$. Therefore, when $\mu_A = \nu_{x_1}$, $\mu(A|x_m, y) < \mu^*(\beta)$ for all $y \in Y$.

Clearly, the experiment G cannot be decisive for priors higher than ν_{x_1} because the decision maker chooses an action a when receiving x_1. Therefore, for G to be decisive for some prior, her right action should be a for all y, which means that no outcome of G can induce the decision maker to choose an action b. Similarly, G can never be decisive for priors below ν_{x_1} since for any $\mu_A < \nu_{x_1}$, $Y_a(\mu_A, x_m) = \emptyset$. □

Proof of Proposition 4.5.1: We use the same notations as in the proof of Lemma 4.5.1.

First, we assume that $\delta = 1$. Choose ν_1.

$$\nu_1(\beta) = \frac{\beta f_B(x_1) g_B(y_1)}{f_A(x_1) g_A(y_1) + \beta f_B(x_1) g_B(y_1)}.$$

This ν_1 satisfies the following equation

$$\nu_1 f_A(x_1) g_A(y_1) = (1 - \nu_1) \beta f_B(x_1) g_B(y_1),$$

which means that if $\mu_A = \nu_1$, the decision maker is indifferent between a and b when she observes x_1, y_1. Choose another belief ν_2.

$$\nu_2(\beta) = \min \left\{ \frac{\beta f_B(x_2) g_B(y_1)}{f_A(x_2) g_A(y_1) + \beta f_B(x_2) g_B(y_1)}, \frac{\beta f_B(x_1) g_B(y_2)}{f_A(x_1) g_A(y_2) + \beta f_B(x_1) g_B(y_2)} \right\}.$$

For $\mu_A \in (\nu_1(\beta), \nu_2(\beta))$, the decision maker has an incentive to choose an experiment at the end of period 1 when she observes x_1 or y_1, and costs c is small.

Let us choose a sequence $\beta_i > 0$ such that $\lim_{i \to \infty} \beta_i = 0$. For each i, we can choose $\nu_i \in (\nu_1(\beta_i), \nu_2(\beta_i))$ and $c_i > 0$ so that the decision maker has an incentive to choose a remaining experiment when she observes x_1 or y_1. Then, for each i, the expected payoffs are independent of sequences because she chooses a only if she observes x_1, y_1. Therefore, she wants to minimize expected costs. Let $z_i(FG)$ and $z_i(GF)$ be expected costs under FG and GF, respectively.

$$\begin{aligned} z_i(FG) &= c_i(\nu_i f_A(x_1) + (1-\nu_i) f_B(x_1)), \\ z_i(GF) &= c_i(\nu_i g_A(y_1) + (1-\nu_i) g_B(y_1)). \end{aligned}$$

Since ν_i converges to zero, $f_B(x_1) \leq g_B(y_1)$ for FG to be better than GF.

Now, let us consider another case. Let

$$\nu_3(\beta) = \max\left\{ \frac{\beta f_B(x_1) g_B(y_l)}{f_A(x_1) g_A(y_l) + \beta f_B(x_1) g_B(y_l)}, \frac{\beta f_B(x_2) g_B(y_{l-1})}{f_A(x_2) g_A(y_{l-1}) + \beta f_B(x_2) g_B(y_{l-1})} \right\}.$$

and

$$\nu_4(\beta) = \frac{\beta f_B(x_2) g_B(y_l)}{f_A(x_2) g_A(y_l) + \beta f_B(x_2) g_B(y_l)}.$$

For the same sequence of β_i's, we can choose $\nu_i' \in (\nu_3(\beta_i), \nu_3(\beta_i))$ and c_i' so that the decision maker has an incentive to choose an experiment when she observes x_2 or y_l. In this case, expected costs, which will be denoted by z_i', are the following.

$$\begin{aligned} z_i'(FG) &= c_i'(\nu_i' f_A(x_2) + (1-\nu_i') f_B(x_2)), \\ z_i'(GF) &= c_i'(\nu_i' g_A(y_l) + (1-\nu_i') g_B(y_l)). \end{aligned}$$

Since ν_i' also converges to zero, it should be the case that $f_B(x_2) \leq g_B(y_l)$. Combining the two conditions we have gotten, we come to the conclusion that

$$1 = f_B(x_1) + f_B(x_2) \leq g_B(y_1) + g_B(y_l) \leq 1.$$

Therefore, $f_B(x_1) = g_B(y_1)$ and $f_B(x_2) = g_B(y_l)$, which means that for all $j = 2, \cdots, j-1$, $g_B(y_j) = 0$. One can prove that $f_A(x_1) = g_A(y_1)$ and $f_A(x_2) = g_A(y_l)$ by choosing a divergent sequence $\beta_i' > 0$. Finally, we do not worry about a discounting factor since it is sufficient to consider fixed costs to prove the proposition. $\square$

Proof of Lemma 4.6.1: Simple calculations will prove the Lemma. $\mu_((x_2, y_1) < \mu_((x_1, y_2)$ only if

$$\frac{Prob(x_2, y_1|A)}{Prob(x_2, y_1|B)} < \frac{Prob(x_1, y_2|A)}{Prob(x_1, y_2|B)}.$$

Plug the parameters into the above inequality.

$$\frac{rq}{r(1-q)} < \frac{(1-r)p(1-q)}{(1-r)(1-p)q}.$$

The assumption means that the above inequality is true. We need to show that $\mu_((x_3, y_1) < \mu_((x_2, y_2)$. A similar calculation shows that under the assumption, $\mu_((x_3, y_1) < \mu_((x_2, y_2)$. $\square$

Proof of Proposition 4.6.2: Suppose not. Since F has finite realizations, H must have finite realizations. Assume that $Z = \{z_1, \cdots, z_m\}$. Without loss of generality, we can assume that

$$\frac{h_A(z_1)}{h_B(z_1)} > \cdots > \frac{h_A(z_m)}{h_B(z_m)}.$$

If F is a combination of G and H, then x_1 must correspond to (y_1, z_1) and x_3 must correspond to (y_2, z_m). This means that

$$(1-r)p = qh_A(z_1),$$
$$(1-r)(1-p) = (1-q)h_A(z_m).$$

From this, we can see

$$\frac{p}{1-p} = \frac{q}{1-q}\frac{h_A(z_1)}{h_A(z_m)}.$$

In addition, (y_1, z_m) and (y_2, z_1) must correspond to x_2.

$$r = qh_A(z_m) = (1-q)h_A(z_1).$$

This means that

$$\frac{h_A(z_1)}{h_A(z_m)} = \frac{q}{1-q}.$$

Then,

$$\frac{p}{1-p} = \frac{q}{1-q}\frac{h_A(z_1)}{h_A(z_m)} = \left(\frac{q}{1-q}\right)^2.$$

This is a contradiction. □

Proof of Lemma 4.6.2: The set of utility vectors are determined by $(1, 0)$, $(0, \beta)$, and the following two vectors:

$$v_1 = ((1-r)p, (r + (1-r)p)\beta),$$
$$v_2 = ((1-r)p + r, (1-r)p\beta).$$

F is more informative than G if and only if $(q, q\beta)$ is a convex combination of v_1 and v_2. Simple algebra shows that $(q, q\beta)$ is a convex combination of v_1 and v_2 if and only if $r \in \left[0, \frac{p-q}{p-\frac{1}{2}}\right]$. □

Proof of Proposition 4.6.1: We consider all cases in terms of μ^* and the posterior beliefs.

Case 1. $\mu(x_1 y_1) \leq \mu^*(\beta)$. In this case, all posterior beliefs are lower than or equal to the threshold belief. The decision maker should run no experiment, and choose action b.

Case 2. $\mu(x_1y_2) < \mu^*(\beta) < \mu(x_1y_1)$**.** The decision maker could have an incentive to run a second experiment only if she receives x_1 or y_1. Consider the strategy under which the decision maker runs G, and take an action based on the realization of G. Since F is more informative than G, the decision maker can do better by running F first and taking an action.

Consider the strategy σ_{GF} under which the decision maker runs G first, and runs F if she observes y_1. This strategy is dominated by the following strategy σ_{FG}: the decision maker runs F first and runs G when x_1 is observed. Under the both strategies, the decision maker chooses a only if she observes x_1 and y_1. This means that under both strategies, the decision maker will choose action a with probability of receiving (x_1, y_1) and choose action b with the complementary probability. Therefore, the expected payoffs are the same, regardless of whether the decision maker follows σ_{GF} or σ_{FG}. What matters is the expected costs.

$$\mathbb{E}[c|\sigma_{GF}] = \mu_0 q + (1-\mu_0)(1-q),$$
$$\mathbb{E}[c|\sigma_{FG}] = \mu_0(1-r)p + (1-\mu_0)(1-r)(1-p).$$

Since $r > \frac{p-q}{p}$, $1 - r < \frac{q}{p}$. The maximum of $\mathbb{E}[c|\sigma_{FG}]$ is

$$\mathbb{E}[c|\sigma_{FG}] \leq \mu_0 q + (1-\mu_0)q\frac{1-p}{p}.$$

Since $p > q$, $q\frac{1-p}{p} < 1 - q$. Therefore, $\mathbb{E}[c|\sigma_{FG}] < \mathbb{E}[c|\sigma_{GF}]$.

Case 3. $\mu(x_2y_1) \leq \mu^*(\beta) \leq \mu(x_1y_2)$**.** The decision maker does not need to run G. Suppose the decision maker observes x_1. The right action is a, and no outcomes of G will change the right action. If the decision maker observes x_2 or x_3, the right action is b, and no outcomes of G will change the right action.

Case 4. $\mu(A|x_2y_2) < \mu^*(\beta) < \mu(A|x_2y_1)$**.** By the same argument as in Case 2, we do not have to consider the case in which the decision maker runs G first, and take the right action without running F.

Suppose the decision maker runs G first, and then runs F based on the realization. If the decision maker runs F in period 2, expected costs are either $(\mu_0 q + (1-\mu_0)(1-q))c$ or $(\mu_0(1-q) + (1-\mu_0)q)c$ or any convex combination of the two. Note that the expected costs cannot be smaller than $(1-q)c$. However, the decision maker can do better by running F first. When the decision maker runs F and observes x_1 or x_3, she does not run G. So, the decision maker has an incentive to run G only if she observes x_2, in which case she incurs the expected costs of rc. Also, under this strategy, the expected payoffs are the largest. Therefore, if the decision maker runs F first, she receives larger expected payoffs and incurs smaller expected costs. Therefore, the decision maker prefers FG to GF.

Case 5. $\mu(A|x_3y_1) \leq \mu^*(\beta) \leq \mu(A|x_2y_2)$. This case is the same as Case 3. The optimal strategy is to conduct F first and take the right action.

Case 6. $\mu(A|x_3y_2) < \mu^*(\beta) < \mu(A|x_3y_1)$**.** This case is similar to Case 2. The decision maker has an incentive to run a second experiment only if she observes x_3 or y_2. If the decision maker chooses GF, the expected costs are

$$\mathbb{E}[c|GF] = \mu_0(1-q) + (1-\mu_0)q.$$

If the decision maker chooses FG, the expected costs are

$$\mathbb{E}[c|FG] = \mu_0(1-r)(1-p) + (1-\mu_0)(1-r)p.$$

As in Case 2, the expected payoffs are the same regardless of the sequence of experiment. Since the expected cost under FG are smaller that what it is under GF, the decision maker prefers FG to GF.

Case 7. $\mu^*(\beta) \leq \mu(A|x_3y_2)$**.** The best choice for the decision maker is to choose action a. She has no incentive to run an experiment.

As we have shown, in every case, the optimal strategy for the decision maker is either to run F first or to take the right action. This completes the proof. $\square$

Proof of Proposition 4.7.1: The structure of the proof is the same as in the proof of Proposition 4.4.3. If $\delta = 1$, the decision maker will run the two experiments, so there exists $\underline{\delta}$ such that for all $\delta \in [\underline{\delta}, 1)$, the decision maker is willing to sequentially run the two experiments if necessary.

We calculate utility vectors. Let σ_{FG} be the optimal strategy in Σ_{FG}.

$$v_A(\sigma_{FG}) = (1-p)u_A(b) + \delta p \left(qu_A(a) + (1-q)u_A(b)\right),$$
$$v_B(\sigma_{FG}) = pu_B(b) + \delta(1-p)\left((1-q)u_B(a) + qu_B(b)\right).$$

Let σ_{GF} be the optimal strategy in Σ_{GF}.

$$v_A(\sigma_{GF}) = (1-q)u_A(b) + \delta q \left(pu_A(a) + (1-p)u_A(b)\right),$$
$$v_B(\sigma_{GF}) = qu_B(b) + \delta(1-q)\left((1-p)u_B(a) + pu_B(b)\right).$$

$\Delta v = v(\sigma_{FG}) - v(\sigma_{GF})$ is the difference between the two utility vectors.

$$\Delta v = (1-\delta)(p-q)(-u_A(b), u_B(b)).$$

If $\Delta v \cdot (\mu_0, 1-\mu_0) \geq 0$, FG is better than GF. Simple calculation yields

$$\mu_0 \leq \frac{u_A(b)}{u_A(b) + u_B(b)}.$$

This completes the proof of the first part. The argument for the second part is symmetric, and hence omitted. $\square$

Proof of Proposition 4.8.1: Choose an increasing sequence $(\beta_n, \mu_{0,n})$ such that the decision maker's current action is b and she changes her action only if she observes x_1 and y_1. Choose an increasing sequence (δ_n) so that the decision maker runs the second experiment if she observes x_1 or y_1.

We first construct the utility vector for sequential experiment FG.

$$v_A^{FG} = (1 - f_A(x_1))u_A(b) + f_A(x_1)\delta(g_A(y_1)u_A(a) + (1 - g_A(y_1))u_A(b)),$$
$$v_B^{FG} = (1 - f_B(x_1))u_B(b) + f_B(x_1)\delta(g_B(y_1)u_B(a) + (1 - g_B(y_1))u_A(b)).$$

For the sequence of experiments GF,

$$v_A^{GF} = (1 - g_A(y_1))u_A(b) + g_A(y_1)\delta(f_A(x_1)u_A(a) + (1 - f_A(x_1))u_A(b)),$$
$$v_B^{GF} = (1 - g_B(y_1))u_B(b) + g_B(y_1)\delta(f_B(x_1)u_B(a) + (1 - f_B(x_1))u_B(b)).$$

Let $\Delta = v^{FG} - v^{GF}$ be the difference between two utility vectors.

$$\Delta_A = u_A(b)\left[(1 - f_A(x_1))(1 - \delta g_A(y_1)) - (1 - g_A(y_1))(1 - \delta f_A(x_1))\right],$$
$$\Delta_B = u_B(b)\left[(1 - f_B(x_1))(1 - \delta g_B(y_1)) - (1 - g_B(y_1))(1 - \delta f_B(x_1))\right]$$

We need both arguments to be non-negative. After manipulating $\Delta_A \geq 0$ and $\Delta_B \geq 0$, we obtain

$$g_A(y_1) \geq f_A(x_1),$$
$$g_B(y_1) \geq f_B(x_1).$$

We can also imagine another situation in which the decision maker's current action is a and she changes her action only if she observes x_2 or y_l.

Similar calculation yields

$$g_A(y_l) \geq f_A(x_2),$$
$$g_B(y_l) \geq f_B(x_2).$$

Therefore, $g_A(y_1) + g_A(y_l) \geq f_A(x_1) + f_A(x_2) = 1$. Also, $g_B(y_1) + g_B(y_l) \geq f_B(x_1) + f_B(x_2) = 1$. This means that for every $\omega \in \{A, B\}$, $g_\omega(y_1) = f_\omega(x_1)$ and $g_\omega(y_l) = f_\omega(x_2)$. For $l' = 2, \cdots, l-1$, $g_\omega(y_{l'}) = 0$. This means that if FG is better than GF, $F = G$. This completes the proof. $\square$

Proof of Proposition 4.9.1: The proof is similar to the proof of Proposition 4.6.1.

Case 1. $\mu(x_1y_1) \leq \mu^*(\beta)$. In this case, all posterior beliefs are lower than or equal to the threshold belief. The decision maker should run no experiment, and should choose action b.

Case 2. $\mu(x_1y_2) < \mu^*(\beta) < \mu(x_1y_1)$**.** The decision maker could have an incentive to run a second experiment only if she receives x_1 or y_1. Consider the strategy under which the decision maker runs G, and take the right action based on the realization of G without running F. Since F is more informative than G, the decision maker can do better by running F first and taking the right action.

Consider the strategy σ^{GF} under which the decision maker runs G first, and runs F if she observes y_1. We compare σ^{GF} with another strategy σ^{FG}, under which the decision maker runs F first, and runs G when x_1 is observed.

We calculate utility vectors. For the sequence GF,

$$\begin{aligned}
v_A^{GF} &= (1-q)u_A(b) + q\delta((1-r)pu_A(a) + (1-(1-r)p)u_A(b)),\\
v_B^{GF} &= qu_B(b) + (1-q)\delta((1-r)(1-p)u_B(a) + (r+(1-r)p)u_B(b)).
\end{aligned}$$

For the sequence FG,

$$\begin{aligned}
v_A^{FG} &= (1-(1-r)p)u_A(b) + (1-r)p\delta(qu_A(a) + (1-q)u_A(b)),\\
v_B^{FG} &= (r+(1-r)p)u_B(b) + (1-r)(1-p)\delta((1-q)u_B(a) + qu_B(b)).
\end{aligned}$$

The difference is:

$$\begin{aligned}
\Delta_A &= u_A(b)\left[(1-p(1-r))(1-\delta q) - (1-q)(1-\delta p(1-r))\right],\\
\Delta_B &= u_B(b)\left[(r+(1-r)p)(1-\delta(1-q)) - q(1-\delta(1-r)(1-p))\right].
\end{aligned}$$

Simplifying the above two equations,

$$\begin{aligned}
\Delta_A &= u_A(b)(1-\delta)(q-(1-r)q),\\
\Delta_B &= u_B(b)(1-\delta)(1-q-(1-r)(1-p)).
\end{aligned}$$

The condition that $r > \frac{p-q}{p}$ means that Δ_A is always non-negative. Since $1-q > 1-p$ and $r < 1$, Δ_B is always non-negative, too. Therefore, σ^{FG} is better than σ^{GF}.

Case 3. $\mu(x_2y_1) \leq \mu^*(\beta) \leq \mu(x_1y_2)$**.** The decision maker does not need to run G. Suppose the decision maker observes x_1. The right action is a, and no outcomes of G will change the right action. If the decision maker observes x_2 or x_3, the right action is b, and no outcomes of G will change the right action.

Case 4. $\mu(A|x_2y_2) < \mu^*(\beta) < \mu(A|x_2y_1)$**.** By the same argument as in Case 2, we do not have to consider the case in which the decision maker runs G first, and take the right action without running F.

Also, we can rule out the case in which the decision maker runs G first, and then runs F with probability one. In this case, the decision maker can be better off if she runs F first, and runs G when receiving x_2. The expected utility is the same under both strategies, however if the decision maker chooses the latter strategy, she has a higher chance of taking the right action earlier.

Only remaining cases are when the decision maker runs G first, and then runs F based on either y_1 or y_2, but not both. For these cases, we will show that the following strategy dominates them. Let σ^{FG} be the strategy under which the decision maker runs F first, and runs G if she observes x_2.

Let us analyze the case when the decision maker runs G first, and runs F when she observes y_1. In this case, the decision maker chooses action a only if she observes (x_1, y_1) or (x_2, y_1). We calculate utility vectors.

$$v_A^{GF} = u_A(a)q\delta\Big[(1-r)p + r\Big] + u_A(b)\Big[1 - q + q\delta(1-r)(1-p)\Big],$$
$$v_B^{GF} = u_B(a)(1-q)\delta\Big[(1-r)(1-p) + r\Big] + u_B(b)\Big[q + (1-q)\delta(1-r)p\Big].$$

When the decision maker follows σ^{FG},

$$v_A^{FG} = u_A(a)\Big[(1-r)p + \delta rq\Big] + u_A(b)\Big[(1-r)(1-p) + \delta r(1-q)\Big],$$
$$v_B^{FG} = u_B(a)\Big[(1-r)(1-p) + \delta r(1-q)\Big] + u_B(b)\Big[(1-r)p + \delta rq\Big].$$

The difference is

$$\Delta_A = u_A(a)(1-r)p(1-\delta q) + u_A(b)\Big[(1-r)(1-p)(1-\delta q) - (1-q)(1-\delta r)\Big],$$
$$\Delta_B = u_B(a)(1-r)(1-p)(1-\delta(1-q)) + u_B(b)\Big[(1-r)p(1-\delta(1-q)) - q(1-\delta r)\Big].$$

Given $u_A(a)$, the infimum of Δ_A occurs when $u_A(b) = u_A(a)$.

$$\begin{aligned}\inf \Delta_A &= u_A(a)\Big[(1-r)(1-\delta q) - (1-q)(1-\delta r)\Big] \\ &= u_A(a)(1-\delta)(q-r) \\ &\geq 0.\end{aligned}$$

Given $u_B(b)$, the infimum of Δ_B occurs when $u_B(a) = u_B(b)$.

$$\begin{aligned}\inf \Delta_B &= u_B(b)\Big[(1-r)(1-\delta(1-q)) - q(1-\delta r)\Big] \\ &= u_A(b)(1-\delta)(1-q-r) \\ &\geq 0.\end{aligned}$$

Therefore, σ^{FG} is better. The analysis for the other case is symmetric, and thus omitted.

Case 5. $\mu(A|x_3y_1) \leq \mu^*(\beta) \leq \mu(A|x_2y_2)$. This case is the same as Case 3. The optimal strategy is to conduct F first and take the right action.

Case 6. $\mu(A|x_3y_2) < \mu^*(\beta) < \mu(A|x_3y_1)$**.** This case is similar to Case 2. The decision maker has an incentive to run a second experiment only if she observes x_3 or y_2. If the decision maker chooses GF,

$$v_A^{GF} = u_A(a)\Big[q + (1-q)\delta(1-(1-r)(1-p))\Big] + u_A(b)(1-q)\delta(1-r)(1-p),$$
$$v_B^{GF} = u_B(a)\Big[1-q+q\delta(1-(1-r)p)\Big] + u_B(b)q\delta(1-r)p,$$

For the sequence FG,

$$v_A^{FG} = u_A(a)\Big[1-(1-r)(1-p)+\delta(1-r)(1-p)q\Big] + u_A(b)\delta(1-r)(1-p)(1-q),$$
$$v_B^{FG} = u_B(a)\Big[1-(1-r)p+\delta(1-r)p(1-q)\Big] + u_B(b)\delta(1-r)pq.$$

The difference is:

$$\Delta_A = u_A(b)\left[(1-(1-r)(1-p))(1-\delta(1-q)) - q(1-\delta(1-r)(1-p))\right],$$
$$\Delta_B = u_B(a)\left[(1-(1-r)p)(1-\delta q) - (1-q)(1-\delta(1-r)p)\right].$$

After some calculation,

$$\Delta_A = u_A(b)(1-\delta)(1-q-(1-r)(1-p)),$$
$$\Delta_B = u_B(b)(1-\delta)(q-(1-r)p).$$

Since $1-q > 1-p$ and $r < 1$, Δ_A is always non-negative. The condition that $r > \frac{p-q}{p}$ means that Δ_B is always non-negative, too. Therefore, σ^{FG} is better than σ^{GF}.

Case 7. $\mu^*(\beta) \leq \mu(A|x_3y_2)$. The best choice for the decision maker is to choose action a. She has no incentive to run an experiment.

So far, we have analyzed the optimal strategy in every binary decision problem, and have shown that the optimal strategy always indicates that the decision maker runs F first. Therefore, FG is better than GF. $\square$

Proof of Proposition 4.9.2: Recall that

$$\begin{aligned}
v_\omega(\sigma_{FG}) &= \sum_{\alpha\in\mathcal{A}} u_\omega(\alpha)\sigma^0_{FG}(\alpha) + \sigma^0_{FG}(F)\left(-c+\delta\sum_{\substack{\alpha\in\mathcal{A},\\ x\in X}} u_\omega(\alpha)\sigma^1_{FG}(\alpha,x)f_\omega(x)\right)\\
&+ \sigma^0_{FG}(F)\delta\sum_{x\in X}\sigma^1_{FG}(G,x)f_\omega(x)\left(-c+\delta\sum_{\substack{\alpha\in\mathcal{A},\\ y\in Y}} u_\omega(\alpha)\sigma^2_{FG}(\alpha,x,y)g_\omega(y)\right)
\end{aligned}$$

We can decompose the expected utility into two terms. The first term is due to discounting and the second term is the expected cost.

$$\begin{aligned}
\zeta_\omega(\sigma_{FG}) = &\sum_{\alpha\in\mathcal{A}} u_\omega(\alpha)\sigma^0_{FG}(\alpha) + \sigma^0_{FG}(F)\delta\sum_{\substack{\alpha\in\mathcal{A},\\ x\in X}} u_\omega(\alpha)\sigma^1_{FG}(\alpha,x)f_\omega(x)\\
&+ \sigma^0_{FG}(F)\sum_{x\in X}\sigma^1_{FG}(G,x)f_\omega(x)\delta^2\sum_{\substack{\alpha\in\mathcal{A},\\ y\in Y}} u_\omega(\alpha)\sigma^2_{FG}(\alpha,x,y)g_\omega(y),\\
\xi_\omega(\sigma_{FG}) = c&\left[\sigma^0_{FG}(F) + \sigma^0_{FG}(F)\delta\sum_{x\in X}\sigma^1_{FG}(G,x)f_\omega(x)\right].
\end{aligned}$$

Let $\vec{\zeta}(\sigma_{FG}) = (\zeta_A(\sigma_{FG}), \zeta_B(\sigma_{FG}))$, $\vec{\xi}(\sigma_{FG}) = (\xi_A(\sigma_{FG}), \xi_B(\sigma_{FG}))$ and $\vec{\mu} = (\mu, 1-\mu)$. Then, the expected utility is:

$$\mathbb{E}[u|\sigma_{FG}] = \vec{\zeta}(\sigma_{FG})\cdot\vec{\mu} - \vec{\xi}(\sigma_{FG})\cdot\vec{\mu}.$$

The proof of Proposition 4.6.1 means that $\vec{\xi}(\sigma_{FG})\cdot\vec{\mu} \le \vec{\xi}(\sigma_{GF})\cdot\vec{\mu}$ given that σ_{FG} is the best among Σ_{FG} and σ_{GF} is the best among Σ_{GF}. Similarly, the proof of Proposition 4.9.1 means that $\vec{\zeta}(\sigma_{FG})\cdot\vec{\mu} \ge \vec{\zeta}(\sigma_{GF})\cdot\vec{\mu}$ given that σ_{FG} is the best among Σ_{FG} and σ_{GF} is the best among Σ_{GF}. Therefore, $\mathbb{E}[u|\sigma_{FG}] \ge \mathbb{E}[u|\sigma_{GF}]$ for every binary decision problem. $\square$

Bibliography

Aghion, P., Bolton, P., Harris, C. and Jullien, B. (1991), "Optimal Learning by Experiments", *The Review of Economic Studies*, Vol. 58, Issue. 4, pp. 621–654.

Amir, R. and Lazzati, N. (2016), "Endogenous Information Acquisition in Bayesian Games with Strategic Complementarities", *Journal of Economic Theory*, Vol. 163, pp. 684–698.

Arrow, K. (1974), The Limits of Organization, Norton, New York, ISBN 0-393-09323-9.

Athey, S. and Levin, J. (2018), "The Value of Information in Monotone Decision Problems", *Research in Economics*, Vol. 72, Issue. 1, pp. 101–116.

Baron, S. and Myerson, R. (1982), "Regulating a Monopoly with Unknown Costs", *Econometrica*, Vol. 50, No. 4, pp. 911–930.

Battigalli, P. *et al.* (2015), "Self-Confirming Equilibrium and Model Uncertainty", *American Economic Review*, Vol. 105, No. 2, pp. 646–677.

Berk, R.H. (1966), "Limiting Behavior of Posterior Distributions When the Model Is Incorrect", *The Annals of Mathematical Statistics*, Vol. 37, No. 1, pp. 51–58.

Billot, A. *et al.* (2000), "Sharing Beliefs: Between Agreeing and Disagreeing", *Econometrica*, Vol. 68, No. 3, pp. 685–694.

Blackwell, D. (1953), "Equivalent Comparison of Experiments", *Annals of Mathematical Statistics*, Vol. 24, Issue 2, pp. 265–272.

Blackwell, D. and Dubins, L. (1962), “Merging of Opinions with Increasing Information”, *The Annals of Mathematical Statistics*, Vol. 33, No. 3, pp. 882–886.

Blume, L. and Easley, D. (2006), “If You're so Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets”, *Econometrica*, Vol. 74, No. 4, pp. 929–966.

Borovička, J. (2011), “Survival and Long-run Dynamics with Heterogeneous Beliefs under Recursive Preferences”, FRB of Chicago Working Paper, No. 2011-06.

Börgers, T., Hernando-Veciana, Á., Krähmer, D. (2013), “When Are Signals Complements or Substitutes?”, *Journal of Economic Theory*, Vol. 148, Issue. 1, pp. 165–195.

Cabrales, A., Gossner, O. and Serrano, R. (2013), “Entropy and the Value of Information for Investors”, *American Economic Review*, Vol. 103, No. 1, pp. 360–377.

Che, Y. and Mierendorff, K. (2017), “Optimal Sequential Decision with Limited Attention”, Working Paper.

Condie, S. (2008), “Living with Ambiguity: Prices and Survival When Investors Have Heterogeneous Preferences for Ambiguity”, *Economic Theory*, Vol. 36, Issue. 1, pp. 81–108.

Da Silva, C.P.S. (2011), “On Asymptotic Behavior of Economies with Complete Markets: the role of ambiguity aversion”, Working paper.

Datta, M., Mirman, L. and Schlee, E. (2002), “Optimal Experimentation in Signal-Dependent Decision Problems”, *International Economic Review*, Vol. 43, No. 2, pp. 577–607.

De Groot, M.H. (1962), “Uncertainty, Information, and Sequential Experiments”, *The Annals of Mathematical Statistics*, Vol. 33, No. 2, pp. 404–419.

Doob, J. L. (1949), Application of the Theory of Martingales, Colloque international Centre nat. Rech. Sci, Paris, pp. 22–28.

Doraszelski, U. and Escobar, J.F. (2010), "A Theory of Regular Markov Perfect Equilibrium in Dynamic Stochastic Games: Genericity, Stability, and Purification", *Theoretical Economics*, Vol. 5, Issue. 3, pp. 369–402.

Dutta, P.K. (1995), "A Folk Theorem for Stochastic Games", *Journal of Economic Theory*, Vol. 66, Issue. 1, pp. 1–32.

Epstein, L.G. and Schneider, M. (2003), "Recursive Multiple Priors", *Journal of Economic Theory*, Vol. 113, Issue. 1, pp. 1–31.

Epstein, L.G. and Schneider, M. (2007), "Learning under Ambiguity", *The Review of Economic Studies*, Vol. 74, No. 4, pp. 1275–1303.

Ganuza, J.J. and Penalva, J. (2010), "Signal Ordering Based on Dispersion and the Supply of Private Information in Auctions", *Econometrica*, Vol. 78, No. 3, pp. 1007–1030.

Ghirardato, P. and Siniscalchi, M. (2016), "Risk Sharing in the Small and in the Large", Working Paper.

Gilboa, I. and Schmeidler, D. (1989), "Maxmin Expected Utility with Non-unique Prior", *Journal of Mathematical Economics*, Vol. 18, Issue. 2, pp. 141–153.

Gneiting, T. and Raftery A.E. (2007), "Strictly Proper Scoring Rules, Prediction, and Estimation", *Journal of the American Statistical Association*, Vol. 102, No. 477, pp. 359–378.

Guesnerie, R. and Laffont, J.J. (1984), "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm", *Journal of Public Economics*, Vol. 25, Issue. 3, pp. 329–369.

Greenshtein, E. (1996), "Comparison of Sequential Experiments", *The Annals of Statistics*, Vol. 24, No. 1, pp. 436–448.

Hossain, T. and Okui, R. (2013), "The Binarized Scoring Rule", *The Review of Economic Studies*, Vol. 80, Issue. 3, pp. 984–1001.

Hörner *et al.* (2011), "Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \to 1$ and A Folk Theorem", *Econometrica*, Vol. 79, No. 4, pp. 1277–1318.

Jullien, B. (2000), "Participation Constraints in Adverse Selection Models", *Journal of Economic Theory*, Vol. 93, Issue. 1, pp. 1–47.

Kalai, E. and Lehrer, E. (1993), "Rational Learning Leads to Nash Equilibrium", *Econometrica*, Vol. 61, No. 5, pp. 1019–1045.

Karni, E. (2009), "A Mechanism for Eliciting Probabilities", *Econometrica*, Vol. 77, No. 2, pp. 603–606.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005), "A Smooth Model of Decision Making under Ambiguity", *Econometrica*, Vol. 73, Issue. 6, pp. 1849–1892.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2009), "Recursive Smooth Ambiguity Preferences", *Journal of Economic Theory*, Vol. 144, Issue. 3, pp. 930–976.

Kocherlakota, N.R. (1996), "Implications of Efficient Risk Sharing without Commitment", *The Review of Economic Studies*, Vol. 63, No. 4, pp. 595–609.

Kogan, L., Ross, S.A., Wang, J., and Westerfield, M.M. (2006), "The Price Impact and Survival of Irrational Traders", *The Journal of Finance*, Vol. LXI, No. 1, pp. 195–229.

Kogan, L., Ross, S.A., Wang, J., and Westerfield, M.M. (2017), "Market Selection", *Journal of Economic Theory*, Vol. 168, pp. 209–236.

Kullback, S. (1959), Information Theory and Statistics, John Wiley & Sons, ISBN 0-8446-5625-9.

Le Cam, L.M. (1953), "On Some Asymptotic Properties of the Maximum Likelihood Estimates and Related Bayes Estimates", *Univ. Calif. Public. Statist.*, Vol. 1, Issue. 11, pp. 277–330.

Lehmann, E.L. (1988), "Comparing Location Experiments", *The Annals of Statistics*, Vol. 16, No. 2, pp. 521–533.

Liang, A., Mu, X., and Syrgkanis, V. (2018), "Optimal and Myopic Information Acquisition", Proceedings of the 2018 ACM Conference on Economics and Computation, pp. 45–46.

Mailath, G. and Samuelson, L. (2006), Repeated Games and Reputations, Oxford University Press, 2006.

Maskin, E. and Riley, J. (1984), "Monopoly with Incomplete Information", *The RAND Journal of Economics*, Vol. 15, No. 2, pp. 171–196.

Mayskaya, T. (2017), "Dynamic Choice of Information Sources", California Institute of Technology Social Science Working Paper.

McCarthy, J. (1956), "Measures of the Value of Information," *Proceedings of the National Academy of Sciences*, Vol. 42, Issue. 9, pp. 654–655.

Oliveira, H. (2018), "Blackwell's Informativeness Theorem Using Diagrams", *Games and Economic Behavior*, Vol. 109, pp. 126–131.

Ozdenoren, E. and Peck, J. (2008), "Ambiguity Aversion, Games against Nature, and Dynamic Consistency", *Game and Economic Behavior*, Vol. 62, Issue. 1, pp. 106–115.

Peters, H. J. M. and Wakker, P. P. (1986), "Convex Functions on Non-Convex Domains", *Economics Letters*, Vol. 22, Issues. 2-3, pp. 251–255.

Qu, X. (2012), "A Mechanism for Eliciting a Probability Distribution", *Economics Letters*, Vol. 115, Issue. 3, pp. 399–400.

Radner, R. (2000), "Costly and Bounded Rationality in Individual and Team Decision-making", *Industrial and Corporate Change*, Vol. 9, Issues. 4, pp. 623–658.

Radner, R. and Stiglitz, J. E. (1984), "A Nonconcavity in the Value of Information", *in* Bayesian Models in Economic Theory (M. Boyer and R.E. Kihlstrom), Elsevier, ch 3, pp. 33–52.

Rigotti, L., Shannon, C., and Strzalecki, T. (2008), "Subjective Beliefs and Ex Ante Trade", *Econometrica*, Vol. 76, No. 5, pp. 1167–1190.

Savage, L.J. (1971), "Elicitation of Personal Probabilities and Expectations", *Journal of the American Statistical Association*, Vol. 66, Issue. 336, pp. 783–801.

Schwartz, L. (1965), "On Bayes Procedures", *Z. Wahrscheinlichkeitstheorie* Vol. 4, Issue. 1, pp. 10–26.

Siniscalchi, M. (2011), "Dynamic Choice under Ambiguity", *Theoretical Economics*, Vol. 6, Issue. 3, pp. 379–421.

Strzalecki, T. and Werner, J. (2008), "Efficient Allocations under Ambiguity", *Journal of Economic Theory*, Vol. 146, Issue. 3, pp. 1173–1194.