Essays on Decision Theory and Information Economics

by

Zichang Wang

Department of Economics
Duke University

Date: ______________________________

Approved:

____
Todd Sarver, Supervisor

____
Philipp Sadowski

____
Arjada Bardhi

____
Curtis Taylor

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2023

Abstract

Essays on Decision Theory and Information Economics

by

Zichang Wang

Department of Economics
Duke University

Date: ____________________
Approved:

____
Todd Sarver, Supervisor

____
Philipp Sadowski

____
Arjada Bardhi

____
Curtis Taylor

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics in the Graduate School of Duke University
2023

Abstract

This dissertation consists of three chapters. Chapter one is a study in decision theory that analyzes regret and information avoidance. Chapter two is a study in information theory that characterizes the comparison of ambiguous information. Chapter three discusses a persuasion model with a constrained sender.

In chapter one, we study regret and information avoidance. Empirical evidence suggests that individuals selectively avoid information, depending on a relevant past choice or lack thereof. We address these findings by studying an agent whose choice behavior can be modeled as if she trades off two conflicting effects of information. The first is a psychological cost from the regret about past choices that are revealed to be suboptimal by the information, whereas the second is the instrumental value of information for making better-informed choices in the future. The primitive of our study is the agent's preference over pairs consisting of a set of menus and an information structure. A set of menus captures a three-period decision problem. Our main axioms reflect the agent's desire to limit her options in period one and to have more flexibility in period two. We posit axioms that connect the agent's consumption choice and information choice. A subjective version of the model is examined where the agent's information choice is not observable. We show that all parameters in both versions of the model can be uniquely identified from the choice behavior.

In chapter two, we study informativeness orders over ambiguous information structures. We generalize Blackwell (1951)'s informativeness order to ambiguous

experiments. The ambiguity in experiments is rooted in a lack of understanding about their probabilistic content. Formally, an ambiguous experiment is modeled as a mapping from an auxiliary state space to the set of unambiguous experiments. We show that one ambiguous experiment is preferred to another by every decision maker for every decision problem if and only if they are related by a condition called *prior-by-prior dominance*, which states that for any first-order belief the decision maker entertains on the auxiliary state space, the expected experiment resulting from this belief for the first experiment is Blackwell more informative than that of the second. This equivalence is robust across a wide range of ambiguity preferences. Comparisons of sets of experiments evaluated using the maxmin criterion are studied as a special case and are shown to result in a weaker informativeness order called *Wald-more-informative*, which states that for any Blackwell experiment in the convex hull of the first set of experiments, there exists another in the convex hull of the second set that is Blackwell less informative.

In chapter three, we study a Bayesian persuasion problem where the persuader's choices of signals are constrained. Specifically, we model this constraint as an α-constraint: Probabilities of any signal realization being sent out conditional on any state of the world are bounded between α and $1-\alpha$. Under this constraint, we extend the "revelation principle" style result in persuasion games by showing that considering the signal realization space to be subsets of the action space is without loss of generality. But it is possible that recommending a proper subset of all actions is uniquely optimal. This possibility contrasts the existing result that having the signal realization space equal to the action space can always be optimal. Based on the revelation principle, we give an algorithm to solve the general constrained persuasion problems. We also provide a characterization of feasible distribution over posterior beliefs for the binary-state-binary-action case, and a comparison of the α-constraint and other existing constraints on the signal space.

Acknowledgements

First and foremost, I wish to thank my advisor, Todd Sarver. I am deeply indebted to Todd for his invaluable advice, guidance, and support for virtually all my research projects. Through his mentorship, Todd has set such high standards for being a scholar and a mentor that I can only hope to meet by myself. Despite the difficulty, I sincerely hope to meet the standards one day in the future.

I also wish to thank Philipp Sadowski for his continued guidance and support. I first became interested in decision theory in one of Philipp's classes. I could not have wished for a better team other than Todd and Philipp to guide me in exploring the realm of decision theory.

I have been very fortunate also to have Arjada Bardhi and Curtis Taylor serve on my dissertation committee. I wish to thank them for their helpful discussions and insightful comments. In particular, I wish to thank Arjada for helping me to navigate the literature on Bayesian persuasion and information design. A special thanks goes to Tracy Lewis for his help and encouragemnt during early stages of my graduate study.

It has been a long journey since the beginning of my study of economics. This journey would not have been possible without the love and support of my parents. This journey has been made so much better with the company of my wife, Kaili. I dedicate this dissertation to them.

Contents

List of Tables

List of Figures

Introduction

The study of information has gain increasingly more attention in the field of modern economic theory. However, simultaneously accumulated is empirical evidence that cannot be satisfactorily explained by standard economic models. This dissertation aims to provide novel theoretical models to explain two lines of empirical findings. The first is individuals' behavior of selective information avoidance, and the second is comparisons of ambiguous information structures.

Information avoidance describes an individual's strict preference for less information even when this information is free and valuable (in the sense that it helps the individual to make better decisions in the future). This behavioral pattern cannot be accommodated in standard decision making models like the expected utility framework. What makes it more puzzling is the subtle observation that this avoidance of information is often selective, depending on a past choice or lack thereof. More specifically, the individual would exhibit information avoidance after making a certain choice, but the same individual would not avoid the same piece of information if there is no relevant choice made in the past. To explain this phenomenon, we develop an axiomatic model in chapter one that uses regret to link new information and old choices. Through the axioms, we characterize the behavioral foundations for selective information avoidance. To the best of our knowledge, this is the first model that can directly explain selective information avoidance. The model can find applications in dynamic choice problems where information gradually flows in. Examples include investment decisions, medical decisions, etc.

A currently standard way to model information is through Blackwell experiments (or statistical experiments). We find this modeling approach a bit oversimplifying as information is pervasively ambiguous in real life, but such ambiguity cannot be directly modeled through precise experiments. Even without any theoretical foundations, intuitive comparisons of ambiguous information structures have been made repeatedly. In chapter two, we provide a theoretical framework to formalize these intuitive comparisons and generalize the celebrated Blackwell's theorem which provides grounds for comparing precise experiments. The theory finds application in preference over different information sources, too.

Chapters one and two focus on the role of information in individual decision-making. Chapter three slightly switches the gear to study the role of information in a specific form of strategic interactions, persuasion games. Empirical evidence suggest that the communication between the sender and the receiver of a persuasion game is often constrained, in the sense that the sender faces constraints on what kind of information structure could be utilized. We incorporate this consideration through an one-parameter constraint into the standard model of persuasion games and study its impact on the implications of the model. We also compare our constraint with existing constraints in the literature.

Chapter 1

Regret and Information Avoidance

1.1 Introduction

Information avoidance is the active avoidance of freely accessible information relevant to decision-making. It is puzzling because standard economic analyses suggest that information may help an individual make better decisions and can also be ignored at no cost. Thus, absent of strategic considerations, information should never be harmful. Nevertheless, there exists overwhelming evidence that information avoidance is widespread.[1]

Some existing theories can rationalize information avoidance in certain specific contexts.[2] However, few of them can directly account for empirical findings that information avoidance is often connected with choices made in the past. Such connections have been extensively documented in a longstanding literature in psychology.[3] The key observation from this literature is that after making a choice, people exhibit

[1] Lerman et al. (1999), Oster, Shoulson, and Dorsey (2013) and Persoskie et al. (2014) are among the large literature documenting individuals avoiding relevant medical information when they are at risk of certain diseases or health conditions. There is also a rapidly growing literature studying investors' aversion to financial information. Examples include Karlsson, Loewenstein, and Seppi (2009), Sicherman et al. (2016) and Hilbert et al. (2022). For a survey on information avoidance that draws from multiple disciplines, see Golman, Hagmann, and Loewenstein (2017).

[2] Among others, there are Caplin and Leahy (2001), Kőszegi (2003), Brunnermeier and Parker (2005), Dillenberger (2010) and Bénabou and Tirole (2011).

[3] This literature dates back to Festinger (1957), which was later followed by Frey and Wicklund (1978), Frey and Rosch (1984), Frey and Stahlberg (1986) and Jonas et al. (2001).

"selective exposure to information" by seeking supportive information and avoiding contradictory information to that choice.[4] In this chapter, we directly address these findings by developing a model that formally links information avoidance to past choices.

How could an economic agent's preference for avoiding information be driven by a choice made in the past? One natural answer is through regret. Information could reveal an agent's past choice to be suboptimal and cause her to experience a sense of regret for not having chosen a different alternative. We refer to this (psychological) effect as the regret cost of information. In addition to the regret cost, our model will also incorporate the instrumental value of information. This instrumental value is derived from the agent's gain from making better-informed choices in the future. In reality, decisions are often made before all relevant information arrives, and it is common for information to reveal suboptimal past choices and to facilitate future decision-making at the same time. For example, a physical examination that helps to screen for potential health problems might also lead to the discovery that an old lifestyle was unhealthy. Inquiries about employment opportunities from different industries for a new job might reveal to a worker that better early-career choices could have been made.

This novel tradeoff between the regret cost of information and the instrumental value of information is at the core of our model. Our model can thus generate both information avoidance and information-seeking behavior, depending on which one of the two effects is dominant in specific decision situations.

To capture this tradeoff between the regret cost and instrumental value of information, we build a model of decision making under uncertainty that involves choices

[4] Examples more directly related to consumer behavior include Ehrlich et al. (1957) and Brock and Balloun (1967), which observed that consumers may have the tendency to avoid information about products they have considered but did not buy or information about risks of products they have purchased.

in three periods. Choices made before the arrival of information open up the possibility of regret, and choices that can be made after the arrival of information are sources for its instrumental value. The uncertainty is captured by a set of objective states of the world and information generally represents what the agent expects to learn about the state of the world during the course of her decision process.

Specifically, we consider an agent whose final choice can be modeled as an act, which is a function that specifies an outcome in each possible state of the world. In period 3 (the final period), the agent chooses an act from a menu (a set of acts) F after the arrival of information. The information is beneficial for this final choice since the agent's choice of act from F can be conditioned on what the information reveals. The menu F, however, was selected by the agent from a set of menus $\mathbb{F}$ in period 2 before the information arrives. Thus, the information could reveal that another menu G that was forgone from $\mathbb{F}$ is actually superior to the chosen menu F. More precisely, if the information reveals that G contains an act g that has higher value than all acts in F, then the agent would regret having chosen F over G. Finally, the agent in our model makes an initial choice in period 1 between different sets of menus, having in mind the subsequent choices of a menu and an act. Therefore, the first dimension of the choice domain in our model corresponds to sets of menus of acts. This modelling approach reflects a common decision procedure in reality in which the agent narrows down her set of options over time before making a final choice. Although it seems to be complicated, it is actually a simplification of standard dynamic decision problems.[5]

There is a second dimension of the choice domain in our model corresponding to information. We interpret it as that the agent has some control over the information that would appear in her three-period decision problem. We formalize this by mod-

[5] For example, in Kreps and Porteus (1978), an agent's choice in each period could determine her consumption in that period and a set of options she can choose from in the next period. Our model simplifies this kind of problem by restricting the consumption to only take place in the final period.

eling information as an *information structure*. Each information structure consists of a set of signal realizations and a collection of conditional probability distributions describing the likelihood for each signal realization to obtain in different states.[6] An agent anticipates that some information will arrive (i.e., some signal realization will be observed) during her decision process and different signal realizations might carry different values to her. Moreover, she can evaluate an information structure from an ex-ante perspective by averaging the values of all possible signal realizations. This second dimension of the choice domain allows us to interpret the agent's behavior as avoiding information, if she picks a less informative information structure, in the sense of the Blackwell order (Blackwell, 1951, 1953), when a more informative one is available.

In sum, we investigate preferences over pairs consisting of a set of menus and an information structure. We show that by simply observing the agent's preference over these pairs, we can determine whether her choices can be modeled *as if* she trades off the regret cost and the instrumental value of the information. Our main result is a representation theorem that features what we refer to as an *informational tradeoff (IT) representation*.

As a building block for the axiomatic characterization of the IT representation, we also consider a subjective informational tradeoff (SIT) representation in which the information structure is a parameter instead of a choice variable. We show that such a representation can be characterized from an agent's preference over sets of menus alone. Moreover, the information structure (as an unobservable parameter) can be elicited from this preference. This characterization is of interest on its own for two reasons. First, conceptually, the dual roles of information in our model do not depend on whether or not the agent has any control over its content. In

[6] This is the (statistical) experiment studied in Blackwell (1951, 1953) and has become a common way to model information.

other words, the agent might still recognize the regret cost and instrumental value of an information structure even if it is exogenously given. Second, from a more pragmatic point of view, the economic modeler may not always be able to observe the information structure anticipated by the agent (even if it is indeed chosen by the agent). It would thus be valuable to be able to elicit the information structure from the agent's consumption behavior.

We now describe the main axioms for our representations.

We have two main axioms for the SIT representation. The first axiom reflects the agent's desire to limit her options for her period-1 choice, since she might experience regret from comparing the menu she has chosen with the counterfactual outcomes represented by the other menus she didn't choose. Formally, suppose F and G are two menus and $\{F\} \succsim \{G\}$. That is, the agent prefers a singleton set containing only menu F to a singleton set containing only menu G. We assume this implies that the agent would (weakly) prefer to choose F over G whenever both F and G are contained in the same set of menus. Our justification for this assumption is from our interpretation that the agent's menu choice from a set of menus is made before any information could arrive. Therefore, the ex-ante comparison between F and G should not depend on whether other menus are present and should be in line with the ranking of the singleton sets $\{F\}$ and $\{G\}$. However, the agent also takes into account the value of each menu after the arrival of the information when evaluating sets of menus. More specifically, adding menu G to a set of menus $\mathbb{F}$ that already contains menu F will always make it (weakly) worse, since G will never be chosen over F from this set of menus but could contain acts that turns out to be better after some signal realization and contribute to the agent's regret. Summarizing the discussion above, the axiom states that if $\{F\} \succsim \{G\}$ and $F \in \mathbb{F}$, then $\mathbb{F} \succsim \mathbb{F} \cup \{G\}$. We refer to this axiom directly as Ex-Ante Regret.[7]

[7] Readers familiar with the literature on regret might already notice the similarity of this axiom

The second main axiom for the SIT representation reflects the agent's preference for a set of menus that allows her to decide later, all else equal. Formally, suppose $\mathbb{F}$ is a set of menus and F, G are two menus. The axiom states that $\mathbb{F} \cup \{F \cup G\} \succsim \mathbb{F} \cup \{F, G\}$. That is, if two sets of menus correspond to the same set of acts that can be ultimately chosen, then the agent would prefer the set that allows her to postpone her decision on which menu to commit to, because doing so would allow her more options to choose from after the arrival of information. We refer to this axiom as Interim Preference for Flexibility.[8]

Other axioms for the characterization of the SIT representation are Weak Order, Continuity, Independence, Finiteness and Domination. These axioms are more standard in the setting of preference over menus and are simply adapted to our setting of preference over sets of menus of acts.

Our axiomatic exercise to characterize the IT representation builds on the axiomatic characterization of the SIT representation. In other words, all axioms used for establishing a SIT representation will also be utilized for the IT representation. Note, however, that these axioms only discipline the agent's choices over sets of menus for some exogenously given information structure. We posit additional axioms that link the agent's consumption choice to her information choice.

1.1.1 Preview of Results

We now describe the functional form identified from our representation theorems.

Let Ω be a finite set of states of the world. The agent's uncertainty about Ω

and the main axiom (dominance) of Sarver (2008). Despite the similarity, the axioms are imposed on different choice domains. Our axiom can be viewed as an adaption from the framework with menus of lotteries to the framework with sets of menus of acts.

[8] Takeoka (2006) posits an axiom similar to ours in a choice domain with menus of menus of lotteries. In the same domain as Takeoka (2006), Kopylov and Noor (2018) and Stovall (2018) each considers an axiom opposite to our axiom, where the agent either prefers to decide earlier or has interim preference for commitment.

is captured by a prior belief $\pi \in \Delta(\Omega)$.[9] An information structure is a mapping $\sigma : \Omega \to \Delta(S)$ where S is a finite set of signal realizations. Let X be a finite set of outcomes. An (Anscombe-Aumann) act is a mapping from Ω to $\Delta(X)$, specifying a lottery over X as the outcome of this act in each state of the world. Let f denote an act. Let F denote a menu of acts and let $\mathbb{F}$ denote a set of such menus. The informational tradeoff (IT) representation for a pair $(\mathbb{F}, \sigma)$ is

$$W(\mathbb{F}, \sigma) := \max_{F \in \mathbb{F}} \sum_{s \in S} \sigma(s) \Big[U(F, \mu_s^\sigma) - R(F, \mathbb{F}, \mu_s^\sigma) \Big] \tag{1.1}$$

where $\sigma(s)$ is the ex-ante probability that signal realization s is generated while μ_s^σ is the Bayesian posterior if realization s is observed.[10] The function $U(F, \mu_s^\sigma)$ captures the material utility of a menu F under posterior μ_s^σ with

$$U(F, \mu_s^\sigma) := \max_{f \in F} \sum_{\omega \in \Omega} \mu_s^\sigma(\omega) u\big(f(\omega)\big) \tag{1.2}$$

where $u : \Delta(X) \to \mathbb{R}$ is an affine function over lotteries that captures the agent's taste over outcomes. Equation (1.2) describes the value of a menu F as the expected value of the best act in it and this contributes to the instrumental value of information.

Finally, $R(F, \mathbb{F}, \mu_s^\sigma)$ captures the agent's regret for having chosen menu F from a set of menus $\mathbb{F}$ at posterior μ_s^σ. Formally,

$$R(F, \mathbb{F}, \mu_s^\sigma) := K \left[\max_{G \in \mathbb{F}} U(G, \mu_s^\sigma) - U(F, \mu_s^\sigma) \right]. \tag{1.3}$$

That is, the agent's regret is proportional to the difference of the material utility for the menu she has chosen and the highest material utility she could have obtained if she had chosen another menu from $\mathbb{F}$. The intensity of regret is capture by the scalar $K \geq 0$.

[9] For any finite set Y, we use $\Delta(Y)$ to denote the probability simplex over Y.

[10] Formally, $\sigma(s) = \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega)$. For any signal realization s with $\sigma(s) > 0$, the posterior $\mu_s^\sigma \in \Delta(\Omega)$ is $\mu_s^\sigma(\omega) = \frac{\pi(\omega)\sigma(s|\omega)}{\sigma(s)}$, and we set the posterior μ_s^σ to be uniform if $\sigma(s) = 0$.

In summary, equation (1.1) specifies that the agent evaluates the value of the pair $(\mathbb{F}, \sigma)$ by the highest expected net value that can be obtained by committing to some menu F in $\mathbb{F}$. The expected net value of a menu F is the weighted average of its net value after each possible signal realization from the information structure σ. This net value is obtained by subtracting the term capturing the regret as defined in equation (1.3) from the material value of the menu as defined in (1.2), reflecting the tradeoff between the regret cost and the instrumental value of information.

In a subjective information tradeoff (SIT) representation, a set of menus of acts $\mathbb{F}$ is evaluated by

$$V(\mathbb{F}) = \max_{F \in \mathbb{F}} \sum_{s \in S} \sigma(s) \Big[U(F, \mu_s^\sigma) - R(F, \mathbb{F}, \mu_s^\sigma) \Big] \tag{1.4}$$

where all terms $\sigma(s)$, μ_s^σ, $U(F, \mu_s^\sigma)$ and $R(F, \mathbb{F}, \mu_s^\sigma)$ are defined in the same way as in the IT representation. Therefore, equations (1.1) and (1.4) describe the same functional form. However, the important difference is that the information structure σ is a choice variable in equation (1.1) while it is a parameter in equation (1.4). We show that we can elicit all parameters in the SIT representation from the preferences over sets of menus, including the information structure.[11]

We now give a numerical example as an illustration for these representations and that they can generate both information-avoiding and information-seeking behavior.

Example 1.1. Consider a student choosing between several colleges. College 1 offers an economics major (Econ) and a computer science major (CS) but does not allow any student to double major. College 2 offers only the economics major and College 3 offers both majors and allow students to double major. Suppose that upon graduation, students majoring in economics can choose between two jobs, banking

[11] Dillenberger, Lleras, Sadowski, and Takeoka (2014) are the first to study the identification of an agent's subjective information from preferences over menus. We use a technique similar to theirs as part of our identification strategy.

(b) and consulting (c), while a computer science major has only one option to work as a software engineer (e). We thus interpret each major as a menu of jobs, that is, Econ $= \{b, c\}$ and CS $= \{e\}$. And we further interpret each college as a set of majors,[12] that is,

$$\mathbb{F}_1 = \{\text{Econ, CS}\} = \Big\{\{b, c\}, \{e\}\Big\}$$

$$\mathbb{F}_2 = \{\text{Econ}\} = \Big\{\{b, c\}\Big\}$$

$$\mathbb{F}_3 = \{\text{Econ, CS, Econ} \cup \text{CS}\} = \Big\{\{b, c\}, \{e\}, \{b, c, e\}\Big\}$$

Suppose the student's choices are based on the career prospect associated with each job, but that there is some uncertainty about these careers. Further suppose that this uncertainty can be represented by a binary state space $\Omega = \{\omega_1, \omega_2\}$ that captures the relevant labor market conditions. Suppose the student's prior belief is such that the two states are equally likely, that is, $\pi(\omega_1) = \pi(\omega_2) = 0.5$. Suppose the state-dependent utility of the career prospect for each job can be summarized by

u	b	c	e
ω_1	110	100	130
ω_2	60	90	50

Finally, suppose $K \geq 0$ and that the student learns the true state upon graduation. That is, after her major choice but before her job choice, she observes a signal realization generated from a fully revealing information structure described by

σ	s_1	s_2
ω_1	1	0
ω_2	0	1

That is, signal realization s_i will be generated with probability 1 contingent on the

[12] We interpret the option of double major to simply mean the student can choose among all three jobs upon graduation. On a different note, $\mathbb{F}_4 = \{\text{Econ} \cup \text{CS}\}$ represents a college that only offers a double major in economics and computer science, which is different from both $\mathbb{F}_1$ and $\mathbb{F}_3$.

state of the world being ω_i. Therefore, observing s_i helps the student to be sure that ω_i is the state of the world.

We first illustrate the SIT representation by computing the value of each college to the student, taking this information structure as fixed. If the student chooses the economics major from $\mathbb{F}_1$, then banking (b) provides the best career prospect if s_1 is observed and consulting (c) is the better choice if s_2 is observed. However, she would feel a sense of regret for majoring in economics if she observes s_1, because she could have $u(e(\omega_1)) = 130$ if she had majored in CS instead of $\max_{f \in \text{Econ}} u(f(\omega_1)) = 110$. Indeed, the material utility (U) and corresponding disutility from regret (R) for choosing each major from $\mathbb{F}_1$ can be summarized by

U	Econ	CS
s_1	110	130
s_2	90	50

R	Econ	CS
s_1	$20K$	0
s_2	0	$40K$

Since the prior is $\pi(\omega_1) = \pi(\omega_2) = 0.5$, each signal realization is generated with probability 0.5. Therefore, the expected net value for the Econ major in $\mathbb{F}_1$ is $0.5 \times (110-20K)+0.5\times(90-0) = 100-10K$. For CS, it is $0.5\times(130-0)+0.5\times(50-40K) = 90-20K$. Since $K \geq 0$, the student would choose Econ over CS in $\mathbb{F}_1$ and the value of College 1 is thus $V(\mathbb{F}_1) = 100 - 10K$.

For College 2, there is only one major, so there is no regret associated with choosing the wrong major.[13] The value for $\mathbb{F}_2$ can be calculated as $V(\mathbb{F}_2) = 0.5 \times (110-0) + 0.5 \times (90-0) = 100$. Note that as long as $K > 0$, that is, as long as the agent is susceptible to feeling regret, then she would choose College 2 over College 1, since by offering the CS major that would never be chosen over the Econ major, College 1 only opens the student up for disutility from regret. This is exactly the motivation for our axiom on Ex-Ante Regret.

[13] It is fair to ask about the possibility for the student to regret her college choice. We believe it is a natural first step for modeling regret by focusing on the regret generated from more recent decisions.

Lastly, through similar calculations, we can see that the student would choose to double major in College 3 by choosing "Econ $\cup$ CS", and the value for the college is $V(\mathbb{F}_3) = 0.5 \times (130 - 0) + 0.5 \times (90 - 0) = 110$. This value is higher than both values of College 1 and College 2, as doing a double major allows the student to best use the information about the state of the world, which also allows her to avoid feeling regret. This is exactly the motivation for our axiom on Interim Preference for Flexibility.

We next illustrate the IT representation and show that an agent represented by the IT representation may avoid information in some scenarios but seek information in other scenarios. In our example, the student might want to learn less about the true state of the world because of the anticipated regret. Formally, suppose the statistical experiment can be parameterized by its precision θ, that is, the information about the state of the world can be represented by

σ_θ	s_1	s_2
ω_1	θ	$1-\theta$
ω_2	$1-\theta$	θ

where $\theta \in [0.5, 1]$. That is, the information is noisy in the sense that even if s_i is observed, the student cannot be sure if ω_i is the true state. But s_i is still indicative about state ω_i since the Bayesian posterior μ_i after observing s_i is $\mu_i(\omega_i) = \theta \geq 0.5$. Suppose the student has chosen College 1 but can control the precision of the information by controlling θ. Our previous calculation shows that $W(\mathbb{F}_1, \sigma_1) = 100 - 10K$. Similar calculations establish that $W(\mathbb{F}_1, \sigma_{0.5}) = 95$. Therefore, as long as $K > 0.5$, the student would prefer a completely noisy information structure, $\theta = 0.5$, to a fully informative information structure, $\theta = 1$. As depicted in Figure 1.1, the optimal precision could depend on the regret intensity K.

More precisely, when $K = 0.4$ (the blue curve), the student strictly prefers the highest precision ($\theta = 1$) than any other precision. When $K = 0.8$ (the red curve),

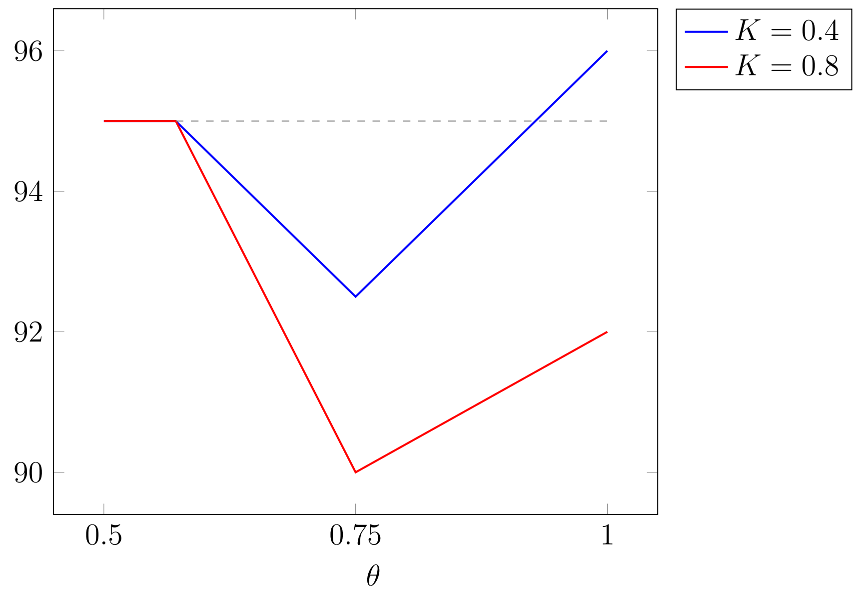

Figure 1.1: Value of College 1, $W(\mathbb{F}_1, \sigma_\theta)$, as a function of θ

sticking with the completely noisy information structure ($\theta = 0$) is optimal.

Note both curves are first flat (for $\theta \in [1/2, 4/7]$), then decreasing (for $\theta \in [4/7, 3/4]$) and then increasing (for $\theta \in [3/4, 1]$). In the flat region, the consulting job (c) will have the best expected payoff after either signal realization s_1 or s_2. The agent thus anticipates no regret by choosing the Econ major. In this region, the value of the college does not depend on the precision of the information structure because it always equals to the ex-ante expected payoff of the consulting job. In the intermediate region ($\theta \in [4/7, 3/4]$), the information is precise enough for the agent to regret choosing the Econ major but not precise enough for her to choose the banking job (b) after observing signal realization s_1. Therefore, the information carries a positive regret cost but zero instrumental value and the increase in its precision only reduces the value of the college. In the final region ($\theta \in [3/4, 1]$), the information is precise enough for the agent to choose banking (b) after observing s_1. The information thus carries a positive instrumental value and the net value of the information starts to increase with its precision.

The rest of the chapter is organized as follows. In Section 1.2, we set up the model and formally describe our primitives. Section 1.3 contains our analysis and results regarding the subjective informational tradeoff representation, with the axioms in Section 1.3.1, the representation theorem in Section 1.3.2, and the uniqueness result in Section 1.3.3. In Section 1.4, we characterize the informational tradeoff representation by considering the larger choice domain where the agent's information choice is also observable. Additional axioms and the representation theorem are presented in Section 1.4.1. Section 1.5 concludes the chapter by discussing some related literature and two extensions.

1.2 The Model

1.2.1 Information Structures

Let Ω be a finite set of states with $|\Omega| \geq 2$. An information structure is a Blackwell experiment with finitely many signal realizations. Formally, an *information structure* is a pair (S, σ) where S is a finite set of signal realizations and σ is a mapping from Ω to $\Delta(S)$. Write $\sigma(s \mid \omega)$ to denote the probability that signal realization s is generated contingent on the state being ω. Different information structures could have different sets of signal realizations. But for convenience, we simply write σ instead of (S, σ). Let $\mathcal{I}$ denote the set of all information structures.

1.2.2 Acts, Menus and Directions

Let X be a finite set of prizes with $|X| \geq 2$ and $\Delta(X)$ is the set of lotteries over X.[14] An (Anscombe-Aumann) *act* is a mapping $f : \Omega \to \Delta(X)$. Let $\mathcal{F}$ be the set of all acts, endowed with the Euclidean metric d. A *menu* is a nonempty compact subset of $\mathcal{F}$, typically denoted by F, G, H. Let $\mathcal{M}$ be the set of all menus. Endow $\mathcal{M}$ with

[14] Equip $\Delta(X)$ with the usual mixture operation and endow $\Delta(X)$ with the Euclidean metric.

the Hausdorff metric d_h.[15] $\mathcal{M}$ is compact.[16] A *direction* is a nonempty compact subset of $\mathcal{M}$, typically denoted by $\mathbb{F}, \mathbb{G}, \mathbb{H}$. A direction is effectively a set of menus. Let $\mathcal{D}$ be the set of all directions. We endow $\mathcal{D}$ with the Hausdorff metric d_H.[17] $\mathcal{D}$ is also compact.

The set $\mathcal{F}$ is equipped with the standard mixture operation. If $f, g \in \mathcal{F}$ and $\alpha \in [0,1]$, then $\alpha f+(1-\alpha)g$ is an act defined by $\big(\alpha f+(1-\alpha)g\big)(\omega) := \alpha f(\omega)+(1-\alpha)g(\omega)$. For any $F, G \in \mathcal{M}$ and $\alpha \in [0,1]$, define the convex combination of these two menus by $\alpha F+(1-\alpha)G := \{\alpha f+(1-\alpha)g \mid f \in F \text{ and } g \in G\}$. Similarly for any directions $\mathbb{F}, \mathbb{G} \in \mathcal{D}$ and $\alpha \in [0,1]$, their convex combination is defined as

$$\alpha\mathbb{F} + (1-\alpha)\mathbb{G} := \{\alpha F + (1-\alpha)G \mid F \in \mathbb{F} \text{ and } G \in \mathbb{G}\}.$$

1.2.3 Primitive

The primitive of our model is a binary relation on $\mathcal{D} \times \mathcal{I}$, representing the agent's preference over pairs consisting of a direction and an information structure. We have in mind an agent facing a three-period decision problem. In period 1, the agent jointly chooses a direction $\mathbb{F}$ and an information structure σ.[18] The agent chooses a menu F from $\mathbb{F}$ in period 2, anticipating the information to arrive after this choice. A signal realization $s \in S$ is then generated according to σ and observed by the agent. In period 3, the agent updates her belief and chooses an act f from F. We

[15] This metric is defined by

$$d_h(F,G) := \max\left\{\max_{f\in F}\min_{g\in G} d(f,g), \max_{f\in G}\min_{g\in F} d(f,g)\right\}.$$

[16] See, for example, Aliprantis and Border (2006, Theorem 3.85).

[17] This is the Hausdorff metric based on d_h, that is, d_H is defined by

$$d_H(\mathbb{F},\mathbb{G}) := \max\left\{\max_{F\in \mathbb{F}}\min_{G\in \mathbb{G}} d_h(F,G), \max_{F\in \mathbb{G}}\min_{G\in \mathbb{F}} d_h(F,G)\right\}.$$

[18] Our model makes no restriction on the order of the direction choice and the information choice. Either choice can be made first, depending on the decision scenario.

do not explicitly model the agent's choices in periods 2 and 3, leaving them as part of the interpretation of the agent's period-1 preference. The timeline is summarized in Figure 1.2.[19]

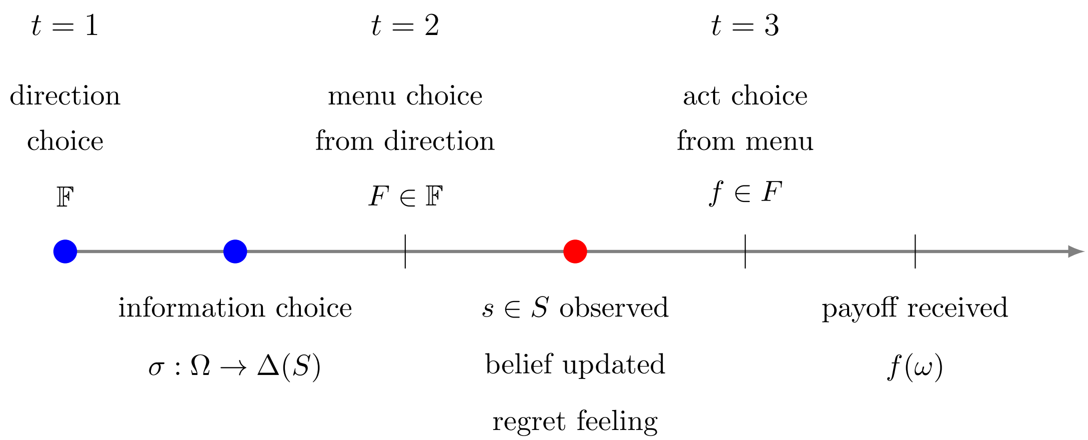

Figure 1.2: Timeline

We start our analysis by restricting attention to a subdomain. Specifically, we consider a binary relation $\succsim$ on $\mathcal{D}$ representing the agent's preference over directions. In this restricted domain, the modeler does not observe the agent's information choice and has to elicit the information structure anticipated by the agent through her preference over directions. The timeline for choices in this subdomain is identical to the timeline described in Figure 1.2. The only difference in interpretation is that the information structure is now an unobservable parameter instead of an observable choice variable. Therefore, we sometimes refer to this as the subjective version of our model. The analysis and results for the subjective version of the model are in Section 1.3. Further building on this analysis and results, the model with the larger choice domain is studied in Section 1.4.

[19] In some applications, it might be more realistic for the information choice to be made after the menu choice. We show that this can be accommodated in a simple extension of our model in Section 1.5.2.

1.3 Subjective Informational Tradeoff Representation

We first give the formal definition of a subjective informational tradeoff representation based on the discussion in the Introduction.

Definition 1.1. A *subjective informational tradeoff (SIT) representation* is a tuple (π, u, K, σ) that consists of a probability measure π on Ω, a non-constant affine function $u : \Delta(X) \to \mathbb{R}$, a constant $K \geq 0$, and an information structure[20] $\sigma : \Omega \to \Delta(S)$ such that $\succsim$ can be represented by the function $V : \mathcal{D} \to \mathbb{R}$ defined by

$$V(\mathbb{F}) = \max_{F \in \mathbb{F}} \sum_{s \in S} \sigma(s) \Big[U(F, \mu_s^\sigma) - R(F, \mathbb{F}, \mu_s^\sigma) \Big] \tag{1.5}$$

where $\sigma(s) = \sum_{\omega \in \Omega} \pi(\omega) \sigma(s \mid \omega)$ is the ex-ante probability of signal realization s, and

- μ_s^σ is the agent's posterior belief after observing s, with

$$\mu_s^\sigma(\omega) = \begin{cases} \frac{\pi(\omega)\sigma(s|\omega)}{\sigma(s)} & \text{if } \sigma(s) > 0 \\ \frac{1}{|\Omega|} & \text{if } \sigma(s) = 0 \end{cases}$$

- $U(F, \mu_s^\sigma)$ is the highest possible expected utility under belief μ_s^σ that can be obtained by choosing some act in menu F. That is,

$$U(F, \mu_s^\sigma) = \max_{f \in F} \sum_{\omega \in \Omega} \mu_s^\sigma(\omega) u\big(f(\omega)\big). \tag{1.6}$$

- $R(F, \mathbb{F}, \mu_s^\sigma)$ is the regret for having chosen F from $\mathbb{F}$ after observing s, that is,

$$R(F, \mathbb{F}, \mu_s^\sigma) = K \left[\max_{G \in \mathbb{F}} U(G, \mu_s^\sigma) - U(F, \mu_s^\sigma) \right], \tag{1.7}$$

  where $K \geq 0$ represents the agent's regret intensity.

[20] We write σ instead of (S, σ) to denote information structures for ease of exposition. As we have mentioned before, we allow different information structures to have different sets of signal realizations.

The interpretation for the SIT representation is just as in the Introduction. When evaluating a direction, the agent anticipates that the information will not arrive until after she makes her menu choice. That is, she will select a menu before observing any signal realization generated from the information structure. On one hand, she can always (weakly) gain from the information by conditioning her choice of act on the signal realizations as captured by equation (1.6) and this constitutes the information value of information. On the other hand, she could experience regret after some signal realizations if her choice of menu is revealed to be inferior as captured by equation (1.7). The regret associated with each signal realization is proportional to the gap between the expected value of the best menu in $\mathbb{F}$ and the expected value of the menu she has chosen. The agent's ex-ante value of the direction under the anticipated information is therefore based on the difference between her expectation of the expected value of the chosen menu and her expectation of the regret.

An useful equivalent expression of the SIT representation can be obtained by combining equations (1.5)-(1.7):

$$\begin{aligned} V(\mathbb{F}) = \max_{F \in \mathbb{F}} & \left[(1+K) \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega) u(f(\omega)) \right] \\ & - K \sum_{s \in S} \max_{G \in \mathbb{F}} \max_{g \in G} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega) u(g(\omega)) \end{aligned} \tag{1.8}$$

Intuitively, the agent in our model chooses a menu from the direction to maximize her expectation of utility minus regret. However, given the information structure, the menu that maximizes her expected utility also minimize her expected regret. That is, the set of maximizers of equation (1.5) within a given direction $\mathbb{F}$ will not depend on the value of K. Therefore, although the regret may cause the agent in our model to prefer smaller directions, it does not distort her choice of menu from a direction since the same set of menus will be the maximizers for the positive term in equation (1.8) no matter how high the regret intensity level is.

1.3.1 Axioms

We impose eight axioms on the binary relation $\succsim$. The first three axioms are standard in the setting of preferences over menus of lotteries and are simply adapted to our setting of preferences over directions (menus of menus) of acts.

Axiom 1.1. (Weak Order): *$\succsim$ is complete and transitive.*

Axiom 1.2. (Continuity): *For any $\mathbb{F}$, the sets $\{\mathbb{G} : \mathbb{F} \succsim \mathbb{G}\}$, $\{\mathbb{G} : \mathbb{G} \succsim \mathbb{F}\}$ are closed.*

Axiom 1.3. (Independence): *For any $\mathbb{F}, \mathbb{G}, \mathbb{H}$ and any $\alpha \in (0,1)$,*

$$\mathbb{F} \succsim \mathbb{G} \iff \alpha\mathbb{F} + (1-\alpha)\mathbb{H} \succsim \alpha\mathbb{G} + (1-\alpha)\mathbb{H}.$$

We refer the reader to Dekel, Lipman, and Rustichini (2001), Dekel, Lipman, Rustichini, and Sarver (2007) and Kopylov (2009) for a discussion of these axioms.

To state the next axiom, we need to introduce the notion of a "critical" subset.

Definition 1.2. We say $\mathbb{G}$ is *critical* for $\mathbb{F}$ if $\mathbb{G} \subseteq \mathbb{F}$ and $\mathbb{H} \sim \mathbb{F}$ for any $\mathbb{H}$ satisfying $\mathbb{G} \subseteq \mathbb{H} \subseteq \mathbb{F}$. We say G is *critical for F in* $\mathbb{F}$ if $G \subseteq F \in \mathbb{F}$ and $(\mathbb{F} \setminus \{F\}) \cup \{H\} \sim \mathbb{F}$ for any menu H satisfying $G \subseteq H \subseteq F$.

The name "critical" is intuitive: If $\mathbb{G}$ is critical for $\mathbb{F}$, then menus in $\mathbb{F}$ but outside $\mathbb{G}$ are all irrelevant for the agent's evaluation of $\mathbb{F}$. The intuition is similar for a menu G being critical for another menu F in a direction $\mathbb{F}$.

Axiom 1.4. (Finiteness): *There exists a natural number N such that:*

- *For every $\mathbb{F}$, there exists $\mathbb{G}$ with $|\mathbb{G}| < N$ such that $\mathbb{G}$ is critical for $\mathbb{F}$.*
- *For every $\mathbb{F}$ and every $F \in \mathbb{F}$, there exists G with $|G| < N$ such that G is critical for F in $\mathbb{F}$.*

In essence, Axiom 1.4 states that only a finite number of menus within any direction matter for evaluating that direction and only a finite number of acts within any menu matter for the evaluation of the direction containing it. Effectively, this restriction reflects the fact the agent is only willing to entertain a finite number of beliefs over Ω as possible posteriors after the information arrival and helps to guarantee that the anticipated information structure has finitely many signal realizations. Axiom 1.4 is adapted from the Finiteness axiom from Stovall (2018), who studies preferences over menus of menus of lotteries. We refer the reader to Stovall (2018) for a more detailed discussion of this axiom and its connection with the finiteness axioms stated in Dekel, Lipman, and Rustichini (2009) and Kopylov (2009).

The following axiom allows for the possibility of regret:

Axiom 1.5. (Ex-Ante Regret): *If* $\{F\} \succsim \{G\}$ *and* $F \in \mathbb{F}$, *then* $\mathbb{F} \succsim \mathbb{F} \cup \{G\}$.

Axiom 1.5 is adapted from the Dominance axiom from Sarver (2008). He studies preferences over menus of lotteries to capture regret generated by an agent's subjective uncertainty about her taste over lotteries. We adapt his dominance axiom to our framework of preferences over directions of acts.

Contrary to standard models, Axiom 1.5 allows for the possibility that $\mathbb{F} \succ \mathbb{F} \cup \{G\}$, that is, the agent may strictly prefer not to add a menu G to a direction $\mathbb{F}$. This reflects the agent's desire to limit the size of the direction in some situations. In addition, Axiom 1.5 specifies that the exact situations in which the agent might exhibit this desire are when the added menu G will never be subsequently chosen from $\mathbb{F} \cup \{G\}$ because an ex-ante better menu F is already contained in $\mathbb{F}$. Intuitively, if G is not chosen in period 2, then it does not benefit the agent to add G to $\mathbb{F}$. But it can harm the agent if some act contained in G turns out to be better after some signal realizations from the anticipated information.

In Example 1.1, the comparison between $\mathbb{F}_1 = \{\text{Econ}, \text{ CS}\}$ and $\mathbb{F}_2 = \{\text{Econ}\}$ is a perfect illustration for these arguments. $\mathbb{F}_1$ can be viewed as adding the option of majoring in CS to $\mathbb{F}_2$, but $\mathbb{F}_1$ is strictly worse than $\mathbb{F}_2$. The reason is that the student still chooses Econ as her major for its better ex-ante value while the option to major in CS makes her regret after realizing that working as a software engineer provides the best career prospect after observing signal realization s_1.

The next axiom reflects the fact that information can still benefit the agent because she values flexibility for her intermediate choice.

Axiom 1.6. (Interim Preference for Flexibility): *For any direction* $\mathbb{F}$ *and menus* F, G*,*

$$\mathbb{F} \cup \{F \cup G\} \succsim \mathbb{F} \cup \{F, G\}.$$

Note that the two directions in the comparison share the same set of acts that can be ultimately chosen. They only differ in how much flexibility can be retained after the arrival of information. Specifically, the agent can choose $F \cup G$ from the direction on the left hand side and choose acts from both F and G after observing any signal realization. But she has to decide between choosing F or G from the direction on the right hand side before the information could arrive. If she selects F, then she has to let go acts that are in G but not in F. Those acts could turn out to be better after certain signal realizations, potentially contributing to both lower payoff in the future and higher regret about the past. Same goes for choosing G. Other than the part about regret, this is the standard argument for the preference for flexibility as coined in Kreps (1979). This axiom differs from the standard argument for the preference for flexibility, however, since our agent not necessarily prefers a direction that contains larger sets.[21]

In Example 1.1, the comparison between $\mathbb{F}_1$ and $\mathbb{F}_3 = \{\text{Econ},\ \text{CS},\ \text{Econ} \cup \text{CS}\}$ perfectly illustrate this axiom. $\mathbb{F}_3$ is strictly preferred to $\mathbb{F}_1$ because the option to

[21] That is, we do not impose any axiom like " $\mathbb{F} \cup \{F\} \succsim \mathbb{F} \cup \{G\}$ whenever $G \subseteq F$."

double major in both Econ and CS in $\mathbb{F}_3$ allows the student to postpone her decision on jobs she has to let go by declaring a single major.

Stovall (2018) considers a related axiom in the setting of preferences over menus of menus of lotteries. His axiom, called "interim preference for commitment," states the opposite of our Axiom 1.6. That is, the agent prefers adding two menus separately to a set of menus comparing to adding their union. In his model, interim preference for commitment is used to capture the fact that the agent is subject to temptation in the intermediate stage.

Our next axiom is the standard nontriviality statement that the agent has strict preference over some pairs of outcomes.

Axiom 1.7. (Nontriviality): *There exist lotteries* $\ell, \ell' \in \Delta(X)$ *with* $\ell \succ \ell'$.[22]

To state the last axiom, we need to introduce the notion of domination.

Definition 1.3. An act f *dominates* another act g if $f(\omega) \succsim g(\omega)$ for all $\omega \in \Omega$.

- Say that a menu F dominates another menu G if for any $g \in G$, there exists $f \in F$ such that f dominates g.
- Say that a direction $\mathbb{F}$ dominates another direction $\mathbb{G}$ if for any $G \in \mathbb{G}$, there exists $F \in \mathbb{F}$ such that F dominates G.

By definition, domination is built on state-by-state comparisons between acts. Therefore, if f dominates g, then the agent will never choose g over f under any belief over Ω. Similarly, if a menu F dominates another menu G, then F will have higher material value than G under any belief. Therefore, there will never be any regret about not choosing G from a direction also containing F. These observations motivate the next axiom.

[22] We follow a common practice of abusing notation by using ℓ to also denote the constant act that yields lottery ℓ in every state. For simplicity, we also write $\ell \succ \ell'$ for $\{\{\ell\}\} \succ \{\{\ell'\}\}$.

Axiom 1.8. (Domination):

1. *If f dominates g, then $f \succsim g$ and $\{\{f, g\}\} \sim \{\{f\}\}$.*

2. *If $\mathbb{F}$ dominates $\mathbb{G}$, then $\mathbb{F} \sim \mathbb{F} \cup \mathbb{G}$.*

The first half of part 1 reflects a sense of monotonicity of the agent's preference. If f is better than g in every state, then f itself should be preferred to g. The second half of part 1 states that adding a dominated act to a menu does not change her attitude toward the menu, because a dominated act contributes to neither the material value of the menu nor the regret. Similarly for part 2, if $\mathbb{F}$ dominates $\mathbb{G}$, then any menu $G \in \mathbb{G}$ is dominated by some menu $F \in \mathbb{F}$. Therefore, adding this collection of dominated menus to $\mathbb{F}$ will not change the agent's choice of menu from $\mathbb{F}$. It will not change the agent's anticipated regret, either. So the agent is indifferent between $\mathbb{F}$ and $\mathbb{F} \cup \mathbb{G}$.

1.3.2 Representation Theorem

Theorem 1.1. *A binary relation $\succsim$ over $\mathcal{D}$ has a subjective informational tradeoff representation if and only if it satisfies Axioms 1.1-1.8.*

The proof of Theorem 1.1 is contained in Appendix A.2. The necessity of the axioms is relatively easy to check.

There are two steps to prove the sufficiency of Axioms 1.1-1.8 for the SIT representation. We first prove a representation theorem (Theorem A.2) for a preference over menus of menus of lotteries that is closely related. For convenience, we refer to a menu of menus of lotteries as a direction of lotteries. The setup for this related choice domain and the results are contained in Appendix A.1. This representation theorem features what we refer to as a "partial regret" representation. It can be viewed as an extended version of the regret representation in Sarver (2008) with three time periods. The agent chooses a direction of lotteries in period 1 and then selects a menu from this direction in period 2. Both choices are made before her subjective uncertainty about her taste over lotteries are resolved. She learns about her taste after the menu choice but before the lottery choice in period 3. This could make her regret her menu choice, but she can also choose the best lottery from this menu according to the revealed taste. Different from Sarver (2008), the agent still has some flexibility after the uncertainty about her taste for lotteries is resolved.

The second step is to establish a connection between the respective choice domains for the SIT representation and the partial regret representation. To do so, we build on a technique used in Dillenberger, Lleras, Sadowski, and Takeoka (2014, henceforth DLST) which involves a sequence of geometric arguments that connects a preference over menus of acts to a preference over menus of lotteries. Our proof connects a preference over directions of acts to a preference over directions of lotteries.

We now apply the SIT representation to two specific types of directions.

Definition 1.4. For any menu F, let $D(F) := \{\{f\} \mid f \in F\}$. We say a direction $\mathbb{F}$ is an *early-commitment direction* if $\mathbb{F} = D(F)$ for some F.

The operation $D(F)$ is one natural way to make a menu F into a direction: We first collect each element $f \in F$ into a singleton menu and then collect all these singleton menus into a direction. In our model, this correspond to the case where the agent has to commit to a final choice (i.e., an act) before she observes any signal realization from the information structure. In other words, the agent is choosing between acts based on her prior. Facing early-commitment directions, information carries zero instrumental value and only contributes to an agent's regret. This is captured by our representation, since when restricting to early-commitment directions,

$$V\big(D(F)\big) = \max_{f \in F} \left[(1+K) \sum_{\omega \in \Omega} \pi(\omega) u(f(\omega)) \right] - K \sum_{s \in S} \max_{g \in F} \sum_{\omega} \pi(\omega) \sigma(s \mid \omega) u(g(\omega)). \tag{1.9}$$

This corresponds to a finite version of the regret representation characterized in Sarver (2008) in the framework with acts. Sarver (2008) studies preferences over menus of lotteries and the agent's regret is from her uncertainty about future tastes for lotteries. Despite the similarity in the way to model regret, his paper does not emphasize the mechanism for the agent's taste change. On the contrary, we focus on the interpretation that information is the driving force for the agent's change in how she evaluates acts.

A direction $\mathbb{F}$ is a *singleton direction* if $\mathbb{F} = \{F\}$ for some menu F. This corresponds to another way of embedding the set of menus into the set of directions. When restricting to singleton directions, the SIT representation reduces to

$$V\big(\{F\}\big) = \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega) \sigma(s \mid \omega) u(f(\omega)). \tag{1.10}$$

That is, a singleton direction corresponds to the case where the agent can make all her relevant decisions after observing a signal realization from the information structure, and this corresponds to a finite version of the subjective learning representation characterized in DLST.

1.3.3 Uniqueness of the SIT Representation

Recall that a subjective informational tradeoff representation for $\succsim$ has four parameters (π, u, K, σ), where π is the prior over Ω, u is an affine function over outcomes, K is a non-negative scalar and σ is an information structure.

Definition 1.5. A *distribution over posteriors,* denoted by ν, is a finitely-supported probability measure on $\Delta(\Omega)$. We say that ν is *induced by* a prior π and an information structure $\sigma : \Omega \to \Delta(S)$ if ν satisfies that for any $\mu \in \Delta(\Omega)$,

$$\nu(\mu) = \sum_{s \in S} \mathbf{1}[\mu = \mu_s^\sigma]\sigma(s),$$

where $\mathbf{1}(\cdot)$ is the indicator function, and $\sigma(s)$ and μ_s^σ are the ex-ante probability of observing s and Bayesian posterior after observing s, respectively.

Since the agent is Bayesian, the induced distribution over posteriors always averages back to the prior, that is, $\sum_{\mu \in \text{supp}(\nu)} \nu(\mu)\mu(\omega) = \pi(\omega)$ for any $\omega \in \Omega$. We say an information structure σ induces a *degenerate* distribution over posteriors if the induced distribution puts weight 1 on the prior belief.

Note from equation (1.5) that if two information structures σ and σ' induce the same distribution over posteriors given a prior π, then the SIT representations with parameters (π, u, K, σ) and (π, u, K, σ') represent the same preference over $\mathcal{D}$.

Theorem 1.2. *Suppose both* $(\pi_0, u_0, K_0, \sigma_0)$ *and* (π, u, K, σ) *represent* $\succsim$*, then*

- $\pi_0 = \pi$.
- $u_0 = au + b$ *for some* $a > 0$ *and* $b \in \mathbb{R}$.
- σ_0 *and* σ *induce the same distribution over posteriors given the prior belief* π.
- $K_0 = K$ *if* σ *induces a nondegenerate distribution over posteriors. If* σ *induces a degenerate distribution over posteriors, then* $K_0, K \in [0, \infty)$.

Proof. See Appendix A.3.1. □

The identifications of the prior belief π and the taste u follow directly from the uniqueness result of the standard subjective expected utility model studied in

Anscombe and Aumann (1963). One quick and intuitive way to understand our uniqueness result on the information structure σ and the regret intensity K is through the uniqueness results from Sarver (2008) and DLST.

When restricting attention to early-commitment directions (i.e., $\mathbb{F} = D(F)$ for some menu F), our representation reduces to a finite version of the regret representation in Sarver (2008) in the context of acts. This is the representation described in equation (1.9). Despite the difference in the primitive, we can think of applying Theorem 4 of Sarver (2008) to jointly identify σ and K as long as σ does not induce a degenerate distribution over posteriors.[23] However, we run into a similar issue as in Sarver (2008) when we try to separately identify σ and K using only the representation in equation (1.9). The issue is that we might not always be able to distinguish between the following two agents. The first is an agent who has a large intensity of regret but anticipates an information structure that is not likely to update her belief away from the prior belief, and the second is an agent who has a low level of regret intensity but anticipates an information structure that is more likely to update her belief away from her prior. More precisely, (σ_1, K) and $(\sigma_2, 2K)$ will generate the same preference over directions according to equation (1.9) if σ_2 is obtained by halving the probability of each signal realization in σ_1 in every state of the world and adding an extra signal realization that is generated with probability 0.5 in every state.

We can overcome this issue and achieve separate identification of the information structure σ and the regret intensity K by taking advantage of our larger choice domain. Specifically, we turn to the singleton directions. When restricting attention to singleton directions (i.e., $\mathbb{F} = \{F\}$ for some menu F), our representation reduces to a finite version of the subjective learning representation in DLST. This is the rep-

[23] When it does, the agent anticipates zero regret from any direction because her evaluations for acts before and after the arrival of information are exactly the same. In such a case, all regret intensity levels will generate the same preference over directions.

resentation described in equation (1.10). DLST show that the information structure (modeled as a distribution over posteriors) can be uniquely identified. This helps us to separately identify σ and K, as long as σ does not induce a degenerate distribution over posteriors.

We can make a sharper statement on the uniqueness of the information structure in the SIT representation, provided that the identified prior belief π has full support. To make the statement, we formally introduce the notions of garbling and Blackwell equivalence.

Definition 1.6. Let $\sigma_1 : \Omega \to \Delta(S_1)$ and $\sigma_2 : \Omega \to \Delta(S_2)$ be two information structures. Say that σ_2 is a *garbling* of σ_1 if there exists $\phi : S_1 \to \Delta(S_2)$ such that

$$\sigma_2(s_2 \mid \omega) = \sum_{s_1 \in S_1} \phi(s_2 \mid s_1)\sigma_1(s_1 \mid \omega)$$

for all $s_2 \in S_2$ and $\omega \in \Omega$.

In words, σ_2 is a garbling of σ_1 if σ_2 can be obtained by adding some noise to the information structure σ_1. Blackwell's theorem establishes that σ_2 is a garbling of σ_1 if and only if σ_1 is *more informative than* σ_2, where "σ_1 being more informative than σ_2" means that every Bayesian expected utility maximizer will weakly prefer σ_1 to σ_2 in every standard decision problem. And we say σ_1 is Blackwell equivalent to σ_2 if σ_1 is more informative than σ_2 and σ_2 is more informative than σ_1.

Corollary. *Suppose both $(\pi_0, u_0, K_0, \sigma_0)$ and (π, u, K, σ) represent $\succsim$. If π has full support, then σ_0 is Blackwell equivalent to σ.*

That is, we can identify the information structure up to its equivalence class specified by the Blackwell informativeness order when the identified prior has full support. The notion of one information structure being more informative than another will also be used to capture the meaning of information avoidance in our model.

1.4 Informational Tradeoff Representation

We now move on to formally establish the informational tradeoff representation. First, we need to come back to the choice domain that involves both a choice of direction and a choice of information structure. Let $\succsim$ be a binary relation over $\mathcal{D} \times \mathcal{I}$,[24] where $\mathcal{D}$ is the set of all directions and $\mathcal{I}$ is the set of all information structures.[25] The interpretation for this binary relation is the same as discussed in Section 1.2.3.

Definition 1.7. A binary relation $\succsim$ over $\mathcal{D} \times \mathcal{I}$ has an *informational tradeoff (IT) representation* if there exists a triple (π, u, K) that consists of a probability measure π on Ω, a non-constant affine function $u : \Delta(X) \to \mathbb{R}$, and a constant $K \geq 0$ such that $\succsim$ can be represented by the function $W : \mathcal{D} \times \mathcal{I} \to \mathbb{R}$ defined by

$$\begin{aligned} W(\mathbb{F}, \sigma) = \max_{F \in \mathbb{F}} & \left[(1 + K) \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega) \sigma(s \mid \omega) u(f(\omega)) \right] \\ & - K \sum_{s \in S} \max_{G \in \mathbb{F}} \max_{g \in G} \sum_{\omega \in \Omega} \pi(\omega) \sigma(s \mid \omega) u(g(\omega)) \end{aligned} \tag{1.11}$$

Equation (1.11) is expressed in the same fashion as equation (1.8): There is no direct reference to the material utility U of a menu or the regret cost R. But one can still see that the first term captures the instrumental value of information since the agent can choose different acts after different signal realizations, and the second term corresponds to the regret cost.

[24] This is an abuse of notation as we also use $\succsim$ to denote a binary relation over $\mathcal{D}$. We will explicitly specify the domain whenever we use $\succsim$ from this point on.

[25] Jakobsen (2021) studies the axiomatic foundation for persuasion models using a related primitive. The modeler in his model observes a sender's preference over information structures indexed by menus of acts and a receiver's choice correspondence from menus indexed by signal realizations. Our works are similar in that both can be used to study preference over information, but our primitive is different in that the modeler in our model observes the agent's preferences over pairs of directions and information structures instead of the conditional preferences along each single dimension of the choice domain.

The IT representation allows richer interpretations of our model comparing to the SIT representation. Fixing an information structure σ, the IT representation fully describes a preference over directions as in the SIT representation. Fixing a direction $\mathbb{F}$, we can now examine the agent's preference over information structures using the IT representation. This enables us to capture the information avoidance behavior of an agent. More precisely, we would say an agent *avoids information* if she exhibits a strict preference for a less informative information structure over an more informative one. The IT representation also allows the comparison of directions across different information structures.

To characterize the IT representation, we first connect a preference over $\mathcal{D} \times \mathcal{I}$ to a collection of preferences over $\mathcal{D}$ indexed by the information structures.

Definition 1.8. Fix a binary relation $\succsim$ over $\mathcal{D} \times \mathcal{I}$ and an information structure $\sigma \in \mathcal{I}$, the *conditional preference* $\succsim^\sigma$ is the binary relation over $\mathcal{D}$ defined by

$$\mathbb{F} \succsim^\sigma \mathbb{G} \text{ if } (\mathbb{F}, \sigma) \succsim (\mathbb{G}, \sigma). \tag{1.12}$$

Through the conditional preferences, we can build on top of the axioms identified for the SIT representation to characterize the IT representation. All our previous axioms (Axioms 1.1-1.8) are imposed on a specific conditional preference and thus have no bite on the agent's choices between directions across different information structures. Our additional axioms will put restrictions that link the agent's direction choice and information choice.

1.4.1 Additional Axioms and Representation Theorem

Let $\succsim$ be a binary relation over $\mathcal{D} \times \mathcal{I}$, and let $\succsim^\sigma$ be the conditional binary relation over $\mathcal{D}$ induced by $\succsim$ and σ. On top of Axioms 1.1-1.8, we impose five additional axioms on $\succsim$ to characterize the IT representation. The first one is the standard rationality requirement for this preference over the larger choice domain.

Axiom 1.9. (Weak Order): $\succsim$ *is complete and transitive.*

Note that if a binary relation $\succsim$ over $\mathcal{D} \times \mathcal{I}$ is complete and transitive, then each conditional preference $\succsim^\sigma$ must also be complete and transitive. Therefore, Axiom 1.9 can be viewed as expanding Axiom 1.1.

The next axiom imposes an independence requirement on mixing directions with acts. For convenience, we write $\alpha\mathbb{F} + (1-\alpha)h$ for $\alpha\mathbb{F} + (1-\alpha)\{\{h\}\}$.

Axiom 1.10. (Act Independence): *For any* $(\mathbb{F}, \sigma)$, $(\mathbb{G}, \sigma')$, *any act* h *and any* $\alpha \in (0,1)$,

$$(\mathbb{F}, \sigma) \succsim (\mathbb{G}, \sigma') \iff (\alpha\mathbb{F} + (1-\alpha)h, \sigma) \succsim (\alpha\mathbb{G} + (1-\alpha)h, \sigma') \tag{1.13}$$

Given the interpretation of our model, the agent's preference over acts (as reflected through the preference for directions like $\{\{f\}\}$) should not depend on the information structure. This is because when the agent has only one act that she can choose, the information can neither help her make better choices in the future nor cause her to regret her choice in the past. Therefore, even though we never formally defined the convex combination of two pairs, we can interpret the pair $(\alpha\mathbb{F} + (1-\alpha)h, \sigma)$ as $\alpha(\mathbb{F}, \sigma) \oplus (1-\alpha)(h, \sigma)$. This could naturally be interpreted as that the agent has chosen σ as her information structure but could either face $\mathbb{F}$ or a single act h. Then the motivation for Axiom 1.10 is the same as the motivation for standard independence axioms.

To state the other axioms, some additional notions are needed. First, an information structure $\sigma \in \mathcal{I}$ is *null* (or *uninformative*) if it always induces a degenerate distribution over posteriors under any prior. It is uninformative in the sense that the agent's posterior belief always equals to the prior facing such an information structure. Let o denote a specific null information structure defined by $o : \Omega \to \Delta(\{s_o\})$. That is, this information structure has only one signal realization s_o and this signal realization is obtained with probability one no matter what the state of the world is. Our next axiom disciplines the agent's behavior when the observed information is exactly o.

Axiom 1.11. (Strategic Rationality when No Information, SRNI): *For any menus* F, G,

$$\{F\} \succsim^o \{G\} \Longrightarrow \{F \cup G\} \sim^o \{F\}.$$

When the agent has chosen the information structure o, there will be no regret. She would evaluate any menu according to the expected value of its best alternative since her choices will be made based on the prior belief. Axiom 1.11 reflects this behavior with the standard strategic rationality statement: If F is preferred to G given o, then F must contain an act that is better than every act in G. Therefore, adding the acts in G to F should not change the agent's attitude toward F.

Suppose $\succsim^o$ has a SIT representation, we will show that Axiom 1.11 (SRNI) guarantees that the information structure identified from $\succsim^o$ agrees with o (i.e., they both induce a degenerate distribution over posteriors). This result will give us an anchor to pin down other information structures.

Any information structure that is Blackwell equivalent to o will always induce a degenerate distribution over posteriors, so o is not the only null information structure. However, we only need to impose the SRNI axiom on the preference over directions conditional on o because the rest of our axioms will imply that Blackwell equivalent information structures induce same conditional preferences over directions.

Let $\sigma : \Omega \to \Delta(S)$ be an information structure and $F \in \mathcal{M}$ be a menu. A *plan* is a mapping $\gamma : S \to F$, with the interpretation that $\gamma(s)$ is the act chosen by the agent if she observes signal realization s. Intuitively, a plan describes the agent's commitment on how to react to the information. For example, if F is a singleton with $F = \{f\}$, then there is only one plan for any S. The plan involves choosing

f after every possible signal realization, because it is the only option. For another example, if $F = \{f, g\}$ and $S = \{s_1, s_2\}$, then the agent has four different plans. The first plan is to choose f after s_1 and choose g after s_2 (we write fg in short). The other three are ff, gf and gg.

Let F^S denote the set of all plans (for a fixed information structure and a menu). Given an information structure $\sigma : \Omega \to \Delta(S)$ and a plan $\gamma \in F^S$, the *act induced by σ and γ*, denoted by γ_σ, is an act defined by

$$\gamma_\sigma(\omega) := \sum_{s \in S} \sigma(s \mid \omega)[\gamma(s)](\omega). \tag{1.14}$$

That is, γ_σ is obtained by reducing the "compound act" described by γ and σ to an act by averaging over different signal realizations. To illustrate, consider again the example where $F = \{f, g\}$ and $S = \{s_1, s_2\}$. The information structure $\sigma : \Omega \to \Delta(S)$ together with the plan fg describe the following compound act: Contingent on ω being the state of the world, the agent will obtain $f(\omega)$ (when the signal realization is s_1 and she chooses f according to the plan fg) with probability $\sigma(s_1 \mid \omega)$ and obtain $g(\omega)$ with probability $\sigma(s_2 \mid \omega)$. The induced act $(fg)_\sigma$ is thus obtained by taking the convex combination $\sigma(s_1 \mid \omega)f(\omega) + \sigma(s_2 \mid \omega)g(\omega)$ state by state, as summarized by equation (1.14).

Definition 1.9. Fix a menu F and some information structure σ, *the menu induced by F and σ* is defined by

$$F_\sigma := \left\{\gamma_\sigma \mid \gamma \in F^S\right\}. \tag{1.15}$$

That is, F_σ is obtained by collecting all the acts that can be induced by the information structure σ and some plan $\gamma \in F^S$.

Axiom 1.12. (Reduction): *For any menu $F \in \mathcal{M}$ and any information structure $\sigma \in \mathcal{I}$,*

$$(\{F\}, \sigma) \sim (\{F_\sigma\}, o).$$

Therefore, the induced menu is $F_\sigma = \{f\}$ for any information structure σ. Axiom 1.12 thus implies that $(f, \sigma) \sim (f, o)$ for any act f and information structure σ.[26] Together with the completeness and transitivity of $\succsim$, Axiom 1.12 implies that the agent has a stable preference over acts that is independent of the information structures. This property of the agent's behavior gives us a common ground to align the conditional preferences indexed by different information structures.

Recall that a early-commitment direction is a direction $D(F) = \{\{f\} \mid f \in F\}$. Our last axiom concerns with mixtures of early-commitment directions $D(F)$ and their corresponding singleton directions $\{F\}$. As it turns out, their interactions have a clear implication on the regret intensity of an agent.

Axiom 1.13. (Balance): *If* $(\{F\}, \sigma_1) \succ (\{F\}, o)$ *and* $(\{G\}, \sigma_2) \succ (\{G\}, o)$*, then for any* $\alpha \in (0, 1]$,

$$\left(\alpha D(F) + (1-\alpha)\{F\}, \sigma_1\right) \sim (\{F\}, o) \iff \left(\alpha D(G) + (1-\alpha)\{G\}, \sigma_2\right) \sim (\{G\}, o).$$

Axiom 1.13 describes the existence of a balance point between the two specific types of directions, the early-commitment directions and the singleton directions. The implicit assumption behind this axiom is as follows. If the agent strictly prefers σ_1 to o when facing a singleton direction $\{F\}$, then the agent will strictly prefer o to σ_1 when facing the corresponding early-commitment direction $D(F)$. The mixture $\alpha D(F) + (1-\alpha)\{F\}$ effectively describes a scenario where the agent needs to first choose an act from F, but has uncertainty about whether she has to stick to this choice (corresponding to $D(F)$) or she has the chance to re-optimize within F (corresponding to $\{F\}$). Axiom 1.13 implies that when facing such an uncertainty, there is a weight α that precisely balances out the expected regret generated from facing $D(F)$ and the expected benefit from facing $\{F\}$. At this balance point, the agent

[26] Similar to what we do in the previous section, we abuse notation a little by writing (f, σ) for the pair $(\{\{f\}\}, \sigma)$.

is indifferent of this option and the option to face $\{F\}$ with no information at all. Moreover, Axiom 1.13 states that this balance point is the same across all menus F and information structures σ as long as the agent strictly benefits from σ compared to no information when facing $\{F\}$.

Axiom 1.13 implies the existence of a unique scalar $\alpha^* \in (0, 1]$ such that

$$\Big(\alpha^* D(F) + (1-\alpha^*)\{F\}, \sigma\Big) \sim (\{F\}, o)$$

for all menu $F \in \mathcal{M}$ and information structure $\sigma \in \mathcal{I}$. Such a unique balance point α^* exists independent of the consumption choice F and the information choice σ because our interpretation for the IT representation involves a single parameter that represents the agent's regret intensity in every decision situation.

We now state our representation theorem for the IT representation.

Theorem 1.3. *A binary relation $\succsim$ over $\mathcal{D} \times \mathcal{I}$ has an informational tradeoff representation if and only if the following conditions are both satisfied:*

- *$\succsim$ satisfies Weak Order, Act Independence, SRNI, Reduction and Balance.*
- *$\succsim^\sigma$ satisfies Axioms 2-8 for each $\sigma \in \mathcal{I}$.*

Moreover, the IT representation is unique as long as the agent has nontrivial preference over information. That is, if both (π_0, u_0, K_0) and (π, u, K) can represent $\succsim$, and $(\mathbb{F}, \sigma) \succ (\mathbb{F}, \sigma')$ for some $\mathbb{F}, \sigma$ and σ', then $\pi_0 = \pi$, $u_0 = au + b$ for some $a > 0$ and $b \in \mathbb{R}$, and $K_0 = K$.

Proof. See Appendix A.3.5. □

1.4.2 Proof Sketch for Theorem 1.3

In this section, we formally state the intermediate results corresponding to the steps hinted in the previous section that help us to build the IT representation.

Definition 1.10. A binary relation $\succsim$ over $\mathcal{D} \times \mathcal{I}$ has an *aligned informational tradeoff representation* if there exists a tuple $\left(\pi, u, (K^\sigma)_{\sigma\in\mathcal{I}}, (i^\sigma)_{\sigma\in\mathcal{I}}\right)$ that consists of a probability measure π on Ω, a non-constant affine function $u : \Delta(X) \to \mathbb{R}$, a collection of information structures $i^\sigma : \Omega \to \Delta(S^\sigma)$ and a collection of non-negative scalars K^σ such that $\succsim$ can be represented by the function $W : \mathcal{D} \times \mathcal{I} \to \mathbb{R}$ defined by

$$\begin{aligned} W(\mathbb{F}, \sigma) = \max_{F\in\mathbb{F}} & \left[(1 + K^\sigma) \sum_{t\in S^\sigma} \max_{f\in F} \sum_{\omega\in\Omega} \pi(\omega) i^\sigma(t \mid \omega) u(f(\omega))\right] \\ & - K^\sigma \sum_{t\in S^\sigma} \max_{G\in\mathbb{F}} \max_{g\in G} \sum_{\omega\in\Omega} \pi(\omega) i^\sigma(t \mid \omega) u(g(\omega)) \end{aligned} \tag{1.16}$$

This is called the "aligned" IT representation because the collection of IT representations for each σ is aligned together by sharing the same prior belief π and the same taste over outcomes u.

Axiom 1.14. (Stable Preference over Acts): *For any f, g and any σ, σ',*

$$(f, \sigma) \succsim (g, \sigma) \iff (f, \sigma') \succsim (g, \sigma')$$

Axiom 1.14 is motivated by the role of information in the decision process. When there is only one act that the agent can choose, any information is irrelevant for the decision since the agent can neither benefit from the information nor be hurt from any regret caused by the information.

Lemma 1.4. *A binary relation $\succsim$ over $\mathcal{D} \times \mathcal{I}$ has an aligned informational tradeoff representation if and only if the following conditions are both satisfied:*

- $\succsim$ *satisfies Weak Order, Stable Preference over Acts and Act Independence.*
- $\succsim^\sigma$ *satisfies Axioms 1.2-1.8 for each $\sigma \in \mathcal{I}$.*

Proof. See Appendix A.3.2. $\square$

That is, by imposing Weak Order, Stable Preference over Acts and Act Independence on top of the axioms that ensure that each conditional preference $\succsim^\sigma$ has a SIT representation, we guarantee the existence of a utility representation for a preference over the two dimensional choice domain and at the same time make sure that the identified prior belief and taste over outcomes are information-independent.

Lemma 1.5. *If $\succsim^o$ has a SIT representation (π^o, u^o, K^o, i^o), then $\succsim^o$ satisfies Axiom 1.11 if and only if i^o induces a degenerate distribution over posteriors.*

Proof. See Appendix A.3.3. □

That is, imposing Axiom 1.11 on an aligned IT representation guarantees that the conditional preference $\succsim^o$ labeled with o indeed corresponds to a situation where the agent is anticipating a null information structure.

Definition 1.11. A binary relation $\succsim$ over $\mathcal{D} \times \mathcal{I}$ has a *regret-varying informational tradeoff representation* if there exists a tuple $\big(\pi, u, (K^\sigma)_{\sigma \in \mathcal{I}}\big)$ that consists of a probability measure π on Ω, a non-constant affine function $u : \Delta(X) \to \mathbb{R}$ and a collection of non-negative scalars K^σ such that $\succsim$ can be represented by the function $W : \mathcal{D} \times \mathcal{I} \to \mathbb{R}$ defined by

$$\begin{aligned} W(\mathbb{F}, \sigma) = \; & \max_{F \in \mathbb{F}} \left[(1 + K^\sigma) \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega) \sigma(s \mid \omega) u(f(\omega)) \right] \\ & - K^\sigma \sum_{s \in S} \max_{G \in \mathbb{F}} \max_{g \in G} \sum_{\omega \in \Omega} \pi(\omega) \sigma(s \mid \omega) u(g(\omega)) \end{aligned} \tag{1.17}$$

Comparing to an aligned IT representation, a regret-varying IT representation has the feature that σ is no longer just a label, it is indeed the information structure anticipated by the agent. Thus, if an agent's preference can be represented by a regret-varying IT representation, then we can model her behavior *as if* she is choosing between information structures taking into account the tradeoff of the benefit for future choices and the regret for past choices. It is not yet the IT representation because the regret intensity level could still vary across different information structures, thus the name "regret-varying."

Lemma 1.6. *A binary relation $\succsim$ over $\mathcal{D} \times \mathcal{I}$ has a regret-varying informational tradeoff representation if and only the following conditions are both satisfied:*

- *$\succsim$ satisfies Weak Order, Act Independence, SRNI and Reduction;*
- *$\succsim^{\sigma}$ satisfies Axioms 1.2-1.8 for each $\sigma \in \mathcal{I}$.*

Proof. See Appendix A.3.4. □

Finally, adding Axiom 1.13 (Balance) on top of a regret-varying IT representation will make sure that the parameters for regret intensity identified from all conditional preferences over directions are the same.

1.5 Discussion and Extensions

We conclude by discussing some related literature and presenting two extensions of our model. In Section 1.5.1, we discuss related work on regret and information avoidance. In Section 1.5.2, we present an extension of our model to allow for the information choice to be made after the menu choice. In Section 1.5.3, we present another extension of our model to discuss the effects of allowing the agent to regret her choice of act.

1.5.1 Related Literature

This chapter contributes to the literature on regret and the literature on information avoidance.

We are not the first to model regret as the difference between what the agent actually gets from a certain choice and the counterfactual best outcome she could have got if she made a different choice (earlier examples include Bell (1982), Loomes and Sugden (1982, 1987) and Sugden (1993)). Sarver (2008) is the first to unveil the possibility of identifying regret through preferences over menus. Through the investigation of preferences over menus of lotteries, Sarver (2008) develops an axiomatic model in which an agent anticipates regret from the resolution of her subjective uncertainty regarding her taste over lotteries. His dominance axiom helps to differentiate regret from other motives for desiring a smaller menu, like temptation. Our work is similar in spirit to capture regret, but we study preferences over a larger choice domain with a emphasis on the role of information. This larger choice domain also helps us to overcome the identification issue for the parameter of regret intensity in Sarver (2008). Buturak and Evren (2017) axiomatize a utility representation similar to Sarver (2008) to explain choice overload where an agent tends to stick more with some default option when there are more options to choose from. Their key assumption is that regret might be asymmetric in the sense that sticking with the default option does not generate regret even if it later turns out to be inferior comparing to some other options, while deviating from the default option opens the agent up to regret. In both papers, the only relevant choice by the agent is made before their subjective uncertainty about their taste is resolved, so this resolution is not beneficial for their future choices. Our model explicitly allows the co-existence of regret for past choices and benefit for future choices. Indeed, this tradeoff between the regret cost of information on past choices and the instrumental value of information through making better-informed future choices is at the core of our analysis. Krähmer and Stone (2013) apply a utility representation similar to the regret representation in Sarver (2008) to argue that ambiguity aversion as captured in Ellsberg's paradox could be explained as an aversion to anticipated regret, because drawing from an urn with an unknown composition of balls opens the agent up to regret while drawing from an unambiguous urn does not generate any regret.

They touch upon the possibility of using regret to explain information avoidance by interpreting ambiguity aversion as a version of information aversion. With a direct focus on explicit information choice, this chapter provides a systematic analysis of this issue. Moreover, our axiomatic approach uncovers the behavioral foundations for regret-driven information avoidance.

We are not the first to attempt to build a theoretical model that can account for information avoidance either. Information avoidance generally has many different appearances in different contexts, and some of these behaviors can be accounted for by existing theories.[27] Caplin and Leahy (2001) introduce a psychological state into the standard expected utility framework and build anticipatory feelings about the future (represented by a belief) into the agent's utility function. This framework, even though non-axiomatic, is very general and allows the possibility for information avoidance driven by different anticipatory feelings, and one of them is anxiety. Kőszegi (2003) utilizes a similar framework with a focus on patient behavior and builds a model where patients avoid relevant medical information in order to avoid the anxiety from the anticipation of the possibility of a bad outcome. Brunnermeier and Parker (2005) model economic agents who can optimally set their own beliefs facing uncertainty, and such an agent might want to avoid information that could break their unwarranted optimistic beliefs.[28] Dillenberger (2010) builds an axiomatic model that features a preference for one-shot over gradual resolution of uncertainty. As a result, the agent in his model might behave as if she is averse to some information if the information represents a partial resolution of uncertainty.

An important feature that distinguishes this chapter from all these works is that we develop a model to formally link information avoidance with past choices. As we

[27] Golman, Hagmann, and Loewenstein (2017) also contains a survey of these theoretical models.

[28] Oster, Shoulson, and Dorsey (2013) uses this theory to reconcile their empirical findings about individuals at risk of having Huntington disease avoiding simple and cheap genetic tests designed to reveal whether they have the disease until after symptoms start to show.

have argued in the Introduction, establishing this link is important because there is a great amount of evidence that people are more likely to avoid information after they make a relevant choice.

Two other relevant papers where past choices could influence preferences are Bénabou and Tirole (2011) and Eyster, Li, and Ridout (2021). Bénabou and Tirole (2011) considers a three-period model that generalizes Caplin and Leahy (2001). Their emphasis is that beliefs are treated as assets by some economic agents because those beliefs are valuable in building an agent's identity. Such agents might dislike information that could contradict the belief they have invested in building. Eyster, Li, and Ridout (2021) builds an axiomatic model to study ex-post rationalization. They study agents who may distort their future choices in order to justify past choices. They thus focus on the specific type of information structure that fully reveals the state of the world after the agent's initial choice. This chapter is different from these two papers in that we build the preference for information directly into the choice domain, which allows us to formally study the connection between an agent's consumption choice and information choice.

1.5.2 Choosing Information after Choosing a Menu

An important feature of the interpretation for the IT representation is that the agent jointly chooses a pair of a direction the information. In particular, as discussed in Section 1.2, the choice of information is made before the choice of menu from the chosen direction. This feature may give the impression that our agent has to commit to a specific information acquisition strategy long before she makes her subsequent choices from a direction, which is somewhat troubling in some potential applications.

To illustrate, consider our Example 1.1 about a student's choice of information about different jobs. In our current interpretation, the student needs to decide on the information structure (i.e., θ, the precision of the noisy signal) and commit to it before she could decide her major. The example would be more convincing if the student's choice about information is made after her major choice.

In this section, we discuss a simple extension of our model that addresses this concern. In this extension, we can interpret the representation as if the agent is choosing an information structure after choosing a menu from the direction. Consider the agent choosing a compact set of information structures, Σ, in period one. That is, Σ is a compact subset of $\mathcal{I}$. This set Σ should be interpreted as the agent's partial commitment about her future information acquisition strategies. Suppose the agent has a preference over pairs of directions and sets of information structures that can be represented by[29]

$$\max_{\sigma\in\Sigma} W(\mathbb{F},\sigma) = \max_{\sigma\in\Sigma}\max_{F\in\mathbb{F}}\sum_{s\in S_\sigma}\sigma(s)\left[\max_{f\in F}\sum_{\omega\in\Omega}\mu_s^\sigma(\omega)u\big(f(\omega)\big) - R(F,\mathbb{F},\mu_s^\sigma)\right] \quad (1.18)$$

$$= \max_{F\in\mathbb{F}}\max_{\sigma\in\Sigma}\sum_{s\in S_\sigma}\sigma(s)\left[\max_{f\in F}\sum_{\omega\in\Omega}\mu_s^\sigma(\omega)u\big(f(\omega)\big) - R(F,\mathbb{F},\mu_s^\sigma)\right] \quad (1.19)$$

where $\sigma(s)$ and μ_s^σ are defined as before, and

$$R(F,\mathbb{F},\mu_s^\sigma) = K\left[\max_{G\in\mathbb{F}}\max_{g\in G}\sum_{\omega\in\Omega}\mu_s^\sigma(\omega)u\big(g(\omega)\big) - \max_{f\in F}\sum_{\omega\in\Omega}\mu_s^\sigma(\omega)u\big(f(\omega)\big)\right].$$

The fact that we can exchange the order of the two maximum operators follows from our agent being Bayesian and forward-looking. And this allows us to interpret this extended model as if the agent chooses her information structure after the menu choice.

[29] We will not present an complete axiomatic treatment for this representation since we are only using this for comparison purposes. But one relatively straightforward way to axiomatize this is to impose the standard strategic rationality axiom on the conditional preference for sets of information structures in addition to the other axioms posited in the main text.

1.5.3 Regretting the Choice of Act

Another important property of the informational tradeoff representation is that the agent only experiences regret once and the regret is only about her choice of menu. This property seems to indicate that the agent in our model cannot regret her choice of act even if she finds out about the true state of the world after she receives her actual payoff.[30] Even though this assumption seems to suggest an inconsistency in our modeling approach, it helps us cleanly isolate the tradeoff of the conflicting effects of information from other potential confounding factors.

In this section, we discuss another simple extension of our model to relax the assumption that the agent cannot regret her choice of act. We'll see that such a relaxation may not result in a change of the agent's preference for information. Therefore, the extended model does not necessarily deliver a different prediction on an agent's information avoidance behavior even though it causes the model to be much more complicated and less intuitive.

We consider an agent who might also regret her choice of act in period 3. This regret could arise because the information arrived before the act choice only partially resolves the uncertainty about the state of the world. Recall that the agent's choice of act is based on a posterior belief, and this posterior belief is not degenerate in general. Therefore, it might be that the agent observes the true state of the world after she receives her payoff, and regret her choice of act if another act might be better given the realized state.[31] To model this, consider an agent who has a preference

[30] Another implicit assumption for the interpretation for the IT representation is that the agent does not regret her direction choice. We believe this assumption may not be as problematic since it is a natural first step to model regret in a way such that more recent decisions are more salient.

[31] Generally, the agent can back out some information about the state of the world by observing the payoff she receives and the partition it induces on the state space Ω. Here we make the simplifying assumption that the agent observes the true state of the world when she receives her payoff.

that can be represented by[32]

$$W_1(\mathbb{F},\sigma) := \max_{F\in\mathbb{F}} \sum_{s\in S} \sigma(s)\Big[\widetilde{U}(F,\mathbb{F},\mu_s^\sigma) - R(F,\mathbb{F},\mu_s^\sigma)\Big] \tag{1.20}$$

where $\sigma(s)$ and μ_s^σ are defined as before. The regret term $R(F,\mathbb{F},\mu_s^\sigma)$ is also defined the same as before, that is,

$$R(F,\mathbb{F},\mu_s^\sigma) = K_0\left[\max_{G\in\mathbb{F}}\max_{g\in G}\sum_{\omega\in\Omega}\mu_s^\sigma(\omega)u\big(g(\omega)\big) - \max_{f\in F}\sum_{\omega\in\Omega}\mu_s^\sigma(\omega)u\big(f(\omega)\big)\right]$$

where $K_0 \geq 0$ represents the regret intensity. However, the "material utility" of a menu F under a posterior belief μ_s^σ, denoted by $\widetilde{U}$, is now also dependent on the direction $\mathbb{F}$ it is chosen from. Formally,

$$\widetilde{U}(F,\mathbb{F},\mu_s^\sigma) := \max_{f\in F}\sum_{\omega\in\Omega}\mu_s^\sigma(\omega)\Big[u\big(f(\omega)\big) - \widetilde{R}(f,F,\mathbb{F},\omega)\Big] \tag{1.21}$$

where $\widetilde{R}(f,F,\mathbb{F},\omega)$ is the regret about the act choice f if the state is revealed to be ω. Formally,

$$\widetilde{R}(f,F,\mathbb{F},\omega) := K_1\left[\max_{G\in\mathbb{F}}\max_{g\in G}u\big(g(\omega)\big) - u\big(f(\omega)\big)\right] + K_2\left[\max_{h\in F}u\big(h(\omega)\big) - u\big(f(\omega)\big)\right] \tag{1.22}$$

where the first term reflects the agent's regret toward the counterfactual outcome she could get if she can re-optimize her choice of menu, while the second term reflects the agent's regret toward the counterfactual outcome she could get if she can only re-optimize her choice of act from the chosen menu F. Parameters $K_1, K_2 \geq 0$ represent the respective regret intensities. Therefore, a representation W_1 in the form of equation (1.20) has five parameters, (π, u, K_0, K_1, K_2).

[32] Again, we will not present an axiomatic treatment for this representation since it is only for comparison purposes.

We are interested in comparing two agents' attitudes toward information. The first one is an agent who can be represented by an IT representation W with parameters (π, u, K), and the second one is an agent who can be represented by W_1 with parameters (π, u, K_0, K_1, K_2). The following lemma states that under some regularity conditions about the regret intensities, the possibility of experiencing about her act choice will not change an agent's preference for information given any direction.

Lemma 1.7. *If representations W and W_1 share the same parameters (π, u) and the regret intensity parameters satisfy $K = \frac{K_0}{1+K_1}$ and $K_2 = 0$, then*

$$W(\mathbb{F}, \sigma) \geq W(\mathbb{F}, \sigma') \iff W_1(\mathbb{F}, \sigma) \geq W_1(\mathbb{F}, \sigma')$$

for any direction $\mathbb{F}$ and information structures $\sigma, \sigma' \in \mathcal{I}$.

Proof. See Appendix A.3.6. □

Intuitively, if $K_2 = 0$ and the agent's regret about her act choice solely comes from the counterfactual comparison with what she could have got from the entire direction $\mathbb{F}$, then this part of regret will only depend on her prior belief but not on her choice of information. And the relative regret intensity levels need to satisfy $K = K_0/(1 + K_1)$. The result hinges on these parametric assumptions. That is, if the relative strength of the regret is mismatched or if the agent's regret about her act choice also comes from the counterfactual comparison with what she could have got from the menu she has chosen, then her preference over information could change. The assumption that $K_2 = 0$ also reflects the subtlety we must face when considering regret in a multiple stage setup: We must be precise about what is the reference point (i.e., the counterfactual outcome) the agent is considering for her comparison.

Another complication is about how anticipatory is the agent toward these regret feelings. For example, in the representation defined in equation (1.20), we assume that the regret term $R(F, \mathbb{F}, \mu_s^\sigma)$ is the same as in the baseline model. That is, we implicitly assume that the agent does not take future regret into account when considering her current regret. It might also be plausible to define the regret term differently by

$$R'(F, \mathbb{F}, \mu_s^\sigma) := K' \left[\max_{G \in \mathbb{F}} \widetilde{U}(G, \mathbb{F}, \mu_s^\sigma) - \widetilde{U}(F, \mathbb{F}, \mu_s^\sigma) \right]$$

where $\widetilde{U}(F, \mathbb{F}, \mu_s^\sigma)$ is defined as in equation (1.21). With this definition, the agent essentially "regret about her future regret" in the sense that she is so forward-looking that she includes her regret from the act choice in the future when evaluating the value of a menu chosen from the direction. Incorporating this consideration would make the model even more complicated. It would be an interesting avenue for future research to find a clean way to incorporate regret into an infinite-horizon discrete time model.

Chapter 2

Informativeness Orders over Ambiguous Experiments

2.1 Introduction

Blackwell (1951, 1953) provides an intuitive way of modeling information by considering information as statistical experiments and establishes an elegant equivalence result on comparisons of such statistical experiments. However, information in reality is pervasively ambiguous, as we rarely know the exact probabilistic content of a statistical experiment when we receive information. For example, a medical test for a certain disease can be thought of as a statistical experiment, and such tests have the possibility of both false positives (the patient does not have the disease but the result is positive) and false negatives (the patient has the disease but the result is negative). Seldom do patients know the exact probabilities of false positive and false negative results, and they may therefore view the distribution of test results as ambiguous.

In this sense, Blackwell's formulation of information as unambiguous statistical experiments seems to be too ideal, and no results were provided on comparing the informativeness of experiments when ambiguity is present, even though some intuitive comparisons can be made in reality (e.g., patients may still intuitively prefer a medical test that has a smaller *maximum* probability of false positive or negative results). To accommodate such intuitive comparisons, can we generalize Blackwell's theorem on the comparison of (unambiguous) experiments to compare the "informativeness" of ambiguous experiments? In other words, can we find a meaningful generalization of Blackwell's *garbling* condition that is equivalent to every decision maker preferring one ambiguous experiment to another in every decision problem?

This chapter provides a positive answer to the question above. Indeed, we can generalize the informativeness notion even when there is ambiguity in the information structures. Moreover, we can show that this new informativeness notion is robust across a wide range of ambiguity preferences.

Let Ω be a finite set of states that are directly relevant to the payoff of the decision maker (henceforth DM). An unambiguous experiment is a mapping $p : \Omega \to \Delta(S)$ where S is a finite set of signal realizations and $\Delta(S)$ is the set of all probability measures over S. The experiment p is unambiguous since the probability that signal s is observed in state ω is unambiguously specified for all s in S and ω in Ω. More than one approach is possible to model ambiguity in experiments. A natural first thought is to consider sets of unambiguous experiments (evaluated using the maxmin criterion), which we analyze in detail in Section 2.4. However, such a modeling approach significantly restricts the richness of ambiguity preferences that can be considered. To allow for more general ambiguity preferences (e.g., the smooth ambiguity preferences), we introduce an auxiliary state space Θ. An ambiguous experiment is modeled as a mapping $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$. For each auxiliary state $\theta \in \Theta$, $\mathbf{p}(\cdot, \theta) : \Omega \to \Delta(S)$ is the corresponding unambiguous experiment. The DM faces ambiguity on Θ. Lack of understanding of the distribution of θ translates to lack of understanding of the probabilistic content of the experiment, making this a convenient (and in some sense canonical) way to model ambiguity about experiments. In particular, as we will show in Section 2.4, this modeling approach nests the comparison of sets of experiments evaluated using the maxmin criterion as a special case. The following two examples further illustrate our modeling approach with an auxiliary state space Θ.

Example 2.1. Consider COVID-19 tests at different hospitals supplied by a single pharmaceutical company. The auxiliary state space Θ could be the set of all possible levels of test precision (in terms of likelihoods of false positives and false negatives). Patients may face ambiguity on Θ since there might be limited data about the precision for a relatively new test (or when applying an existing test to a new variant of the virus). The accuracy of any test clearly depends on Θ, but it can still vary across different hospitals due to different implementations (e.g., some hospitals may have more experienced testing crews than others).

Example 2.2. Consider financial analysis about a firm's financial status conducted by different financial service companies. The auxiliary state space Θ could be a set capturing the aspects that are common to all analysis but open to interpretation, such as the quality of public data or the representativeness (or predictive power) of past observations for future performance. Outside investors may face ambiguity on Θ due to their lack of relevant knowledge. The informational quality of any analysis clearly depends on Θ, but it can still vary across different companies.

We view the auxiliary state space as a modeling device that represents the common source of ambiguity for the ambiguous experiments to be compared. We treat this space and the dependence of the experiments on it as commonly understood by both the modeler and the DM. It is "auxiliary" since it is not directly payoff-relevant and influences the DM's decision and payoff only through influencing the information content of the experiment. Patients' payoffs only directly depend on whether or not they are infected, but not directly on the precision of the test. Investors' profits only directly depend on the financial status of the company, but not directly on the quality of the financial analysis.

2.1.1 Preview of Main Results

Let $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ be an ambiguous experiment. For a probability measure μ over the auxiliary state space Θ, the *expected experiment* $\mathbf{p}_\mu : \Omega \to \Delta(S)$ is defined by taking the expectation of $\mathbf{p}$ with respect to μ, that is, $\mathbf{p}_\mu := \int_\Theta \mathbf{p}(\cdot, \theta)\ d\mu(\theta)$. Expected experiments are unambiguous by construction. Two ambiguous experiments $\mathbf{p}$ and $\mathbf{p}'$ are related by the *prior-by-prior dominance* condition if the expected experiment $\mathbf{p}_\mu$ is Blackwell more informative than $\mathbf{p}'_\mu$ for every μ. We show in Section 2.3 that $\mathbf{p}$ is preferred to $\mathbf{p}'$ (i.e., guarantees a higher ex-ante utility) by every decision maker in every decision problem *if and only if* $\mathbf{p}$ prior-by-prior dominates $\mathbf{p}'$. Prior-by-prior dominance is a direct generalization of Blackwell's informativeness notion: When Θ is a singleton, ambiguous experiments reduce to unambiguous experiments and prior-by-prior dominance reduces to being Blackwell more informative.

If a DM's ambiguity preference could be summarized by a single subjective belief μ over Θ, then the problem of comparing ambiguous experiments $\mathbf{p}$ and $\mathbf{p}'$ can be simplified as comparing their corresponding expected experiments $\mathbf{p}_\mu$ and $\mathbf{p}'_\mu$. This simplification is not possible in general since a DM facing ambiguity may have more general ambiguity preferences (e.g., the multiple prior preferences or the smooth preferences). Nonetheless, our main result implies that if $\mathbf{p}_\mu$ is Blackwell more informative than $\mathbf{p}'_\mu$ for *every* possible probability measure μ over Θ, then the DM will prefer $\mathbf{p}$ to $\mathbf{p}'$ as long as his ambiguity preference can be represented by some monotone aggregator of the auxiliary states, even if this aggregator does not correspond to a single subjective belief. This class of monotone aggregators is extremely general and includes essentially all ambiguity preference models used in the literature.

On the one hand, the prior-by-prior dominance condition is powerful since it is sufficient for guaranteeing higher ex-ante utility across a wide range of ambiguity preferences. On the other hand, it is not overly restrictive, in the sense that it is necessary for guaranteeing higher ex-ante utility within the small class of subjective expected utilities. Therefore, prior-by-prior dominance is a robust equivalence condition for guaranteeing higher ex-ante utility for any class of monotone ambiguity preferences that nests subjective expected utility.

We now give a numerical example to illustrate our comparison results:

Example 2.3. Consider two hospitals offering COVID-19 tests. $\mathbf{p}_1$ summarizes the test offered at Hospital 1 while $\mathbf{p}_2$ summarizes that at Hospital 2. Individuals are either infected (I) or not infected (N), and this infection status is the payoff relevant state. The outcome of a test is either positive ($+$) or negative ($-$). Probabilities of false positives and false negatives are ambiguous, modeled through an auxiliary state space $\Theta = [0, 0.02] \times [0, 0.02]$. A typical element $\theta = (\theta_+, \theta_-)$ denotes that the probability of a false positive is θ_+ and the probability of a false negative is θ_-.

In this example, $\Omega = \{I, N\}$, $S_1 = S_2 = \{+, -\}$, and $\Theta = [0, 0.02] \times [0, 0.02]$. Suppose $\mathbf{p}_1$ and $\mathbf{p}_2$ are given by

$\mathbf{p}_1(\cdot, \theta_+, \theta_-) =$

	$+$	$-$
I	$1-\theta_-$	θ_-
N	θ_+	$1-\theta_+$

$\mathbf{p}_2(\cdot, \theta_+, \theta_-) =$

	$+$	$-$
I	$1-1.01\theta_-$	$1.01\theta_-$
N	$1.01\theta_+$	$1-1.01\theta_+$

where $(\theta_+, \theta_-) \in \Theta$. That is, the test offered at Hospital 2 is more likely to generate false positives and false negatives in every possible state.[1] It can be checked that $\mathbf{p}_1$ prior-by-prior dominates $\mathbf{p}_2$, coinciding with our intuition that $\mathbf{p}_1$ should be regarded as "more informative." To further illustrate, consider a third test $\mathbf{p}_3$ defined by

$\mathbf{p}_3(\cdot, \theta_+, \theta_-) =$

	$+$	$-$
I	$1-\alpha$	α
N	α	$1-\alpha$

[1] One possible justification for the perfect correlation of the tests on Θ is as follows: Both hospitals utilize testing kits from the same pharmaceutical company but Hospital 2 has a less experienced testing crew who increases the probabilities of false positive/negative due to human error.

That is, $\mathbf{p}_3$ is an unambiguous test with a known false positive/negative rate of α. When $\alpha = 0$, $\mathbf{p}_3$ prior-by-prior dominates both $\mathbf{p}_1$ and $\mathbf{p}_2$, and will be considered superior by every decision maker. When $\alpha \geq 0.02$, $\mathbf{p}_1$ prior-by-prior dominates $\mathbf{p}_3$ and is thus preferred. When $\alpha \in (0, 0.02)$, $\mathbf{p}_3$ is in general not comparable with $\mathbf{p}_1$ or $\mathbf{p}_2$, as the decision maker's subjective belief over Θ becomes relevant. For example, if a decision maker believes that both θ_+ and θ_- do not exceed α with probability 1 in a certain decision scenario, then he would prefer $\mathbf{p}_1$ to $\mathbf{p}_3$ in that scenario.

Another possible approach to model the ambiguity in experiments is to consider sets of unambiguous experiments. We show in Section 2.4 that we can treat comparisons of sets of unambiguous experiments as a special case of comparisons of ambiguous experiments by focusing on a specific class of DMs: The DMs who apply Wald's maximin criterion when aggregating payoffs across θ. In this special case, one set of unambiguous experiments P is preferred to another P' by every decision maker who applies the maximin criterion if and only if a *Wald-more informative* condition (W-more informative, in short) holds: We say P is W-more informative than P' if for any unambiguous experiment p in the convex hull of P, there exists p' in the convex hull of P' such that p is Blackwell more informative than p'. We show that the W-more informative condition induces an informativeness order that is weaker than that induced by prior-by-prior dominance, that is, there exists pairs of ambiguous experiments that are not comparable according to the prior-by-prior dominance condition although their corresponding sets of experiments are comparable according to the W-more informative condition. The intuition about why we get a weaker notion of informativeness is as follows: A DM who applies the maximin criterion reduces to an expected utility maximizers only when the set of experiments is a singleton. Thus, this specific class of DMs does not nest the class of DMs with subjective expected utility, which causes the prior-by-prior dominance condition to be not necessary.

For a more concrete illustration, consider Example 2.3 again and define

$$P_1 := \{\mathbf{p}_1(\cdot, \theta) \mid \theta \in \Theta\}, \quad P_3 := \{\mathbf{p}_3(\cdot, \theta) \mid \theta \in \Theta\}.$$

Then in cases where $\alpha = 0$ or $\alpha \geq 0.02$, W-informativeness and prior-by-prior dominance agree with each other: P_3 is W-more informative than P_1 when $\alpha = 0$ and the rank is simultaneously reversed when $\alpha \geq 0.02$. However, when $\alpha \in (0, 0.02)$, P_3 is W-more informative than P_1 even though their corresponding ambiguous experiments are not comparable according to the prior-by-prior dominance condition.

Our results hinge on several assumptions. First, we assume that the DM either has commitment power or is dynamically consistent. This assumption is needed as dynamic inconsistency may arise for some belief updating protocols when ambiguity is present. Second, we assume that the ambiguity is only on the auxiliary state space Θ while the DM has an unambiguous prior over the payoff-relevant state space Ω.[2] Third, we assume that the auxiliary states only have an indirect impact on the DM's payoff. A detailed discussion of the effects of relaxing the second and the third assumptions is in Section 2.5.

2.1.2 Related Literature

We contribute to the literature on the interaction of Blackwell's informativeness order and ambiguity preferences. Models based on the combination of ambiguous priors and unambiguous experiments have been previously considered in the literature. Çelen (2012) shows that the Blackwell informativeness order remains valid among the class of maxmin expected utility preferences, under the assumption that the DM has commitment power. Further expanding the class of ambiguity preferences, Li and Zhou (2016) keep the commitment assumption and validate that the

[2] Like in Example 2.3, the ambiguity is about the precision of the test, but not about the prior probability that a test taker is infected.

Blackwell informativeness order remains valid when the DM possess the very general uncertainty averse preferences as in Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011). In contrast to these studies, we consider a different problem and our study leads to a very different result: In our model, the ambiguity is in the information structures instead of in the prior beliefs over the payoff relevant state space, and our prior-by-prior dominance condition is a non-trivial generalization of the Blackwell informativeness order. Although different from existing papers, the study of ambiguous information structures could be important since information structures are indeed pervasively ambiguous in reality, and ambiguous information has gained increased attention in both the applied literature[3] and the experimental literature[4] as its importance has become recognized. We also discuss in Section 2.5 the combination of ambiguous priors and ambiguous experiments as an extension to our model. Chen (2022) shares our approach of modeling ambiguous experiments as a mapping from an auxiliary state space[5] to the space of Blackwell experiments, but his focus is on the learning behavior under ambiguous experiments, while our focus is on comparing their informativeness.

The closest work to ours is Gensbittel, Renou, and Tomala (2015) (henceforth GRT). In their paper, ambiguous experiments are modeled as convex and compact sets of *joint distributions* over payoff-relevant states and signal realizations. They consider DMs who apply Wald's maximin criterion and study models with and without the commitment assumption. Their case of commitment has a large overlap with the special case of our model discussed in Section 2.4. One major methodological difference sets us apart, both in terms of our results, and in terms of their interpretations. We model unambiguous experiments as collections of conditional distributions

[3] For example, Beauchêne, Li, and Li (2019), Bose and Renou (2014), Chen (2022), and Epstein and Schneider (2007, 2008).

[4] For example, Epstein and Halevy (2019), Liang (2019) and Shishkin and Ortoleva (2019).

[5] An auxiliary state space is referred to as "a model space" in his paper.

over signal realizations given the state, while they model them as joint distributions over states and signal realizations. Although our results are superficially similar, they are not directly comparable unless extra assumptions regarding the payoff-relevant prior are imposed. By their modeling approach, GRT treat the DM's prior belief over Ω as part of the information structure instead of part of the DM's preference parameters. Therefore, they are interpreting the theory as if the modeler can observe the DM's payoff-relevant priors. In contrast, our modeling approach treats the DM's prior belief over Ω as a preference parameter. Therefore, our comparison is prior-free in the sense that it can be applied when the modeler cannot observe DM's payoff-relevant priors.[6] In addition, GRT study a model where the commitment assumption is relaxed and the DMs still apply the maximin criterion, while we impose the commitment assumption but consider a much more general class of ambiguity preferences.

The rest of the chapter is organized as follows. We review Blackwell's theorem in Section 2.2. General comparisons of ambiguous experiments are studied in Section 2.3. The special case of comparisons of sets of experiments is studied in Section 2.4. Variations of the key assumptions and their impact on our results are considered in Section 2.5. Proofs omitted in the main text are relegated to Appendix B.

2.2 Blackwell's Theorem

In this section, we describe some primitives for comparisons of experiments and review Blackwell's theorem. Let Ω be a finite set of states of the world. For any finite set X, we use $\Delta(X)$ to denote the set of all probability measures over X.

Definition 2.1. A *Blackwell experiment* (or interchageably, an *experiment*, or an *unambiguous experiment*) is a mapping $p : \Omega \to \Delta(S)$ where S is a finite set of signal

[6] See Section 2.4.3 for additional discussion and examples.

realizations and p maps each state $\omega \in \Omega$ to a probability measure over S.

Slightly abusing notation, we will write $p(s \mid \omega)$ to denote the probability of observing s when the state is ω. An experiment can thus be viewed as a $|\Omega| \times |S|$ real matrix with each row representing a probability distribution over S. To define the notion of garbling, we first define the composition of two stochastic operators.

Suppose X, Y, Z are finite sets, and $\alpha : X \to \Delta(Y)$ and $\beta : Y \to \Delta(Z)$ are two stochastic operators. Their *composition* $\beta \circ \alpha : X \to \Delta(Z)$ is defined by

$$(\beta \circ \alpha)(z \mid x) = \sum_{y \in Y} \alpha(y \mid x)\beta(z \mid y), \ \forall (x, z) \in X \times Z.$$

That is, the composition of β and α gives a probability of z given x when the stochastic operator α is applied followed by β. Now we can define the notion of garbling.

Definition 2.2. Given two Blackwell experiments $p : \Omega \to \Delta(S)$ and $p' : \Omega \to \Delta(S')$, p' is a *garbling* of p if there exists some $\gamma : S \to \Delta(S')$ such that $p' = \gamma \circ p$.

Intuitively, p' is a garbling of p if one can replicate p' (in terms of the conditional probability of s' given ω) by adding "noise" (applying a stochastic operator) to p.

Consider an individual with a finite set of actions A facing a set S of signal realizations (we assume S to be finite throughout the chapter). An **action plan** is a mapping $\sigma : S \to \Delta(A)$. We write $\sigma(\cdot \mid s)$ to denote the individual's (mixed) strategy after observing signal realization s. We use A_S to denote the collection of all action plans once the set of actions A and the set of signal realizations S are fixed, that is, $A_S := \{\sigma \mid \sigma : S \to \Delta(A)\}$.

Consider a Bayesian expected-utility maximizer with a finite set of actions A, a state-dependent utility function $u : \Omega \times A \to \mathbb{R}$, and a prior $\pi \in \Delta(\Omega)$. For this individual, the expected utility from action plan σ for experiment p is

$$U(\sigma, p) := \sum_{s \in S} \sum_{\omega \in \Omega} \sum_{a \in A} \pi(\omega) p(s \mid \omega) \sigma(a \mid s) u(\omega, a) \tag{2.1}$$

Definition 2.3. Facing A, u and π, the individual's **ex-ante expected utility** from the experiment p is $\max_{\sigma \in A_S} U(\sigma, p)$.

We now review Blackwell's theorem.

Theorem 2.1 (Blackwell (1951, 1953))**.** *Given two Blackwell experiments $p : \Omega \to \Delta(S)$ and $p' : \Omega \to \Delta(S')$, the following are equivalent:*

1. *p' is a garbling of p.*
2. *For any A, u, π and any action plan $\sigma' \in A_{S'}$, there exists an action plan $\sigma \in A_S$ such that $U\big(\sigma, p\big) = U\big(\sigma', p'\big)$.*
3. *Every Bayesian expected utility maximizer prefers p to p' for any possible decision problem. That is, p gives weakly higher ex-ante expected utility than p' for every A, u, and π.*

For a simple and elegant proof of Blackwell's theorem, see de Oliveira (2018). If any of the conditions in Theorem 2.1 holds, we say p is **(Blackwell) more informative** than p', and write $p \trianglerighteq p'$. Blackwell's theorem establishes the equivalence of a statistical condition (garbling) and an economical condition (higher ex-ante expected utility for any decision problem). Our goal is to study ambiguous experiments and characterize the equivalent condition (generalization of "garbling") to all decision makers with more general ambiguity preferences having higher ex-ante utility.

2.3 Comparing Ambiguous Experiments

As before, let Ω be a finite set of states that are directly relevant to the DM's payoff. Let Θ be a set of auxiliary states that govern the realization of Blackwell experiments. As illustrated in the introduction, Θ represents the source of ambiguity for the experiment. For ease of exposition, we focus on the case where Θ is a finite set in the main text. Our characterization result remains valid for any nonempty Θ. The more general result is stated in Appendix B.1 and proved in Appendix B.2.

Definition 2.4. An *ambiguous experiment* is a mapping $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ where S is a finite set of signal realizations.

For each auxiliary state $\theta \in \Theta$, $\mathbf{p}(\cdot, \theta)$ can be thought of as a mapping from Ω to $\Delta(S)$, that is, $\mathbf{p}(\cdot, \theta)$ is the Blackwell experiment associated with state θ, and a natural interpretation for $\mathbf{p}$ is a mapping from the auxiliary state space to the set of Blackwell experiments. This structure helps to capture the DM's lack of understanding of the probabilistic content of the experiment.

Consider an individual with a finite set of actions A, a state-dependent utility function $u : \Omega \times A \to \mathbb{R}$ and a prior belief $\pi \in \Delta(\Omega)$. For ambiguous experiment $\mathbf{p}$ and action plan σ, the **expected utility conditional on** state θ is

$$U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) = \sum_{s \in S} \sum_{\omega \in \Omega} \sum_{a \in A} \pi(\omega) \mathbf{p}(s \mid \omega, \theta) \sigma(a \mid s) u(\omega, a). \tag{2.2}$$

To clarify, we use $\mathbf{p}(s \mid \omega, \theta)$ to denote the probability of observing s when the pair of states is (ω, θ). This is another slight abuse of notation as we also use $\mathbf{p}(\cdot, \theta)$ to denote the Blackwell experiment associated with auxiliary state θ. These conditional expected utilities will be aggregated through an aggregator over Θ.

Definition 2.5. Say $V : \mathbb{R}^{\Theta} \to \mathbb{R}$ is a *monotone aggregator* if V is continuous[7] and $V(f) \geq V(g)$ whenever two functions $f, g : \Theta \to \mathbb{R}$ satisfy $f(\theta) \geq g(\theta)$ for all $\theta \in \Theta$.

Essentially all ambiguity preferences used in the literature correspond to some monotone aggregator, and special cases of V will be discussed in the next section

[7] Since Θ is finite, we can endow $\mathbb{R}^{\Theta}$ with the Euclidean topology and the continuity is with respect to this standard topology.

to illustrate its generality. Since monotonicity is the only restriction we have on V, flexible ambiguity attitudes (ambiguity averse, ambiguity loving, mixed attitude) are allowed for essentially all classes of ambiguity preferences it nests.

Definition 2.6. Let $V : \mathbb{R}^\Theta \to \mathbb{R}$ be a monotone aggregator that captures the individuals' ambiguity preferences. Given A, u, π, and V, the individual's *ex-ante utility* from the ambiguous experiment $\mathbf{p}$ is

$$\max_{\sigma \in A_S} V\left(\left\{U\big(\sigma, \mathbf{p}(\cdot, \theta)\big)\right\}_{\theta\in\Theta}\right) \tag{2.3}$$

where $U\big(\sigma, \mathbf{p}(\cdot, \theta)\big)$ is the conditional expected utility defined in equation (2.2).

Two implicit assumptions are behind the definition above. First, we assume that the DM's has an unambiguous prior belief π over the payoff-relevant space Ω. Moreover, the payoffs $u(\omega, a)$ are first aggregated into the expected utility conditional on each auxiliary state θ, and the DM's ex-ante utility is the aggregation of these conditional expected utilities by V. Whether or not the DM has a subjective belief regarding the auxiliary state space Θ is generally irrelevant given the utility specified in equation (2.3).[8] Second, we assume that the auxiliary states do not directly affect the payoff of the DM. This reflects our interpretation that the auxiliary states only directly affect the information content of an experiment, but do not affect outcomes of the DM's actions.[9]

The timing of the events is as follows: First, the DM makes an action plan $\sigma \in A_S$ and we assume the DM is dynamically consistent.[10] Second, an auxiliary

[8] We believe this is a reasonable first step to formalize the evaluation of ambiguous information. We will discuss the impact of allowing prior ambiguity on Ω and more general procedures to aggregate $u(\omega, a)$ into an ex-ante utility function in detail in Section 2.5.

[9] See Section 2.1 for more examples. The alternative modeling approach is to consider a mapping $u : \Omega \times \Theta \times A \to \mathbb{R}$ in which the auxiliary states have a direct impact on the DM's payoff for each action. This approach will be discussed in detail in Section 2.5.

[10] As is well known, dynamic inconsistency may arise with belief updating when ambiguity is present. For an example that illustrates the prevalence of violation of dynamic consistency in general dynamic ambiguity preference models, see Example 2 of Asano and Kojima (2019). Alternatively, we can assume the DM can commit to any action plan he makes.

state θ is drawn from Θ but *not* observed by the DM and signal realizations will be generated according to the Blackwell experiment $\mathbf{p}(\cdot, \theta)$. Third, a payoff relevant state ω is drawn from Ω and a signal realization s is drawn from S according to the distribution $\mathbf{p}(\cdot \mid \omega, \theta)$. Last, the DM observes the signal realization s and acts according to his action plan $\sigma(\cdot \mid s)$.

2.3.1 Classes of Ambiguity Preferences

Any decision maker will be identified with his ambiguity preference, which is fully captured by its corresponding monotone aggregator V. Let $\mathcal{V}_{Mono}$ denote the class of monotone aggregators, that is, $\mathcal{V}_{Mono} := \{V : \mathbb{R}^\Theta \to \mathbb{R} \mid V \text{ is monotone}\}$. $\mathcal{V}_{Mono}$ will be the largest (in terms of set inclusion) class of aggregators we study.

Some special cases of V are listed below to illustrate the generality of our approach to represent ambiguity preferences using monotone aggregators.

Subjective expected utility. V is specified by a single prior $\mu \in \Delta(\Theta)$, which represents the DM's subjective belief over Θ. Fixing μ, $V_\mu : \mathbb{R}^\Theta \to \mathbb{R}$ is defined by

$$V_\mu(f) := \int_\Theta f(\theta)\ d\mu(\theta) = \sum_{\theta \in \Theta} f(\theta)\mu(\theta), \tag{2.4}$$

and as in equation (2.3), we apply this aggregator to $f(\theta) = U\big(\sigma, \mathbf{p}(\cdot, \theta)\big)$.[11] Let $\mathcal{V}_{EU}$ denote the class of subjective expected utility aggregators, that is, $\mathcal{V}_{EU} := \{V_\mu \mid \mu \in \Delta(\Theta)\}$.

The multiple prior preferences, introduced in Gilboa and Schmeidler (1989). V is specified by a closed set of priors $M \subset \Delta(\Theta)$, which represents the set of priors the DM is willing to entertain. Fixing M, V_M is defined by

$$V_M(f) := \min_{\mu \in M} \int_\Theta f(\theta)\ d\mu(\theta) = \min_{\mu \in M} \sum_{\theta \in \Theta} f(\theta)\mu(\theta). \tag{2.5}$$

[11] By applying V_μ to equation (2.3), the DM behaves as if he has a belief over $\Omega \times \Theta$ that is independent across Ω and Θ. This sense of "independence" does not carry over to the other listed ambiguity preferences, since they cannot be summarized by a single belief over Θ.

In the representation of Gilboa and Schmeidler (1989), M is required to be closed and convex. We only require the closedness of M to guarantee that the minimum is well-defined. Let $\mathcal{V}_{MP}$ denote the class of aggregators corresponding to the multiple prior preferences, that is, $\mathcal{V}_{MP} := \{V_M \mid M \subset \Delta(\Theta),\ M \text{ closed}\}$.

Wald's maximin criterion. A subclass of aggregators within $\mathcal{V}_{MP}$ can be described by the maximin criterion introduced in Wald (1950), where the DM is willing to entertain all degenerate beliefs δ_θ for each $\theta \in \Theta$. The aggregator V_W is defined by

$$V_W(f) := \min_{\mu \in \{\delta_\theta \mid \theta \in \Theta\}} \sum_{\theta \in \Theta} f(\theta)\mu(\theta) = \min_{\theta \in \Theta} \ f(\theta). \tag{2.6}$$

We will sometimes refer to V_W as the Wald aggregator, and unlike the other classes of aggregators considered in this section, $\{V_W\}$ is not a superset of $\mathcal{V}_{EU}$. We will analyze this subclass in detail in Section 4.

The smooth ambiguity preferences, introduced in Klibanoff, Marinacci, and Mukerji (2005). V is specified by a function $\phi : \mathbb{R} \to \mathbb{R}$ and a second-order belief $\nu \in \Delta(\Delta(\Theta))$.[12] Fixing ϕ and ν, $V_{\phi,\nu}$ is defined by

$$V_{\phi,\nu}(f) := \int_{\Delta(\Theta)} \phi\left(\int_\Theta f(\theta)\ d\mu(\theta)\right) d\nu(\mu) \tag{2.7}$$

In the representation of Klibanoff, Marinacci, and Mukerji (2005), ϕ is required to be strictly increasing and weakly concave, where the concavity captures the DM's aversion to ambiguity. We only require ϕ to be strictly increasing to guarantee that $V_{\phi,\nu}$ is monotone and mixed ambiguity attitudes could be allowed. Let $\mathcal{V}_S$ denote the class of smooth ambiguity preference aggregators, that is, $\mathcal{V}_S := \{V_{\phi,\nu} \mid \phi \text{ is strictly increasing},\ \nu \in \Delta(\Delta(\Theta))\}$.

All classes of ambiguity preferences listed above are special cases of the uncertainty averse preferences introduced in Cerreia-Vioglio, Maccheroni, Marinacci, and

[12] $\Delta(\Delta(\Theta))$ is the set of all probability measures on the Borel σ-algebra of $\Delta(\Theta)$ under the Euclidean topology.

Montrucchio (2011), which also corresponds to a monotone aggregator over the conditional expected utilities.[13]

2.3.2 Prior-by-Prior Dominance

The last piece of the puzzle for establishing the informativeness order over ambiguous experiment is a statistical condition that generalize the garbling condition in Blackwell's theorem. To get to that condition, it is useful to first define expected experiments.

Definition 2.7. Fixing an ambiguous experiment $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ and a probability measure $\mu \in \Delta(\Theta)$, the *expected experiment* with respect to μ , denoted by $\mathbf{p}_\mu$, is a Blackwell experiment defined by $\mathbf{p}_\mu := \sum_{\theta \in \Theta} \mathbf{p}(\cdot, \theta)\mu(\theta)$, or more explicitly

$$\mathbf{p}_\mu(s \mid \omega) := \sum_{\theta \in \Theta} \mathbf{p}(s \mid \omega, \theta)\mu(\theta), \ \forall (s, \omega) \in S \times \Omega. \tag{2.8}$$

Now we can formally state our statistical condition.

Definition 2.8. Let $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ and $\mathbf{p}' : \Omega \times \Theta \to \Delta(S')$ be two ambiguous experiments. We say $\mathbf{p}$ *prior-by-prior dominates* $\mathbf{p}'$ if $\mathbf{p}_\mu$ is Blackwell more informative than $\mathbf{p}'_\mu$ for all $\mu \in \Delta(\Theta)$, and write $\mathbf{p} \trianglerighteq_{PBP} \mathbf{p}'$.

In other words, $\mathbf{p} \trianglerighteq_{PBP} \mathbf{p}'$ if there exists a family of garblings $\{\gamma_\mu\}$ such that

$$\mathbf{p}'_\mu = \gamma_\mu \circ \mathbf{p}_\mu \text{ for all } \mu \in \Delta(\Theta). \tag{2.9}$$

The "prior-by-prior" quantifier in the name of the condition does not refer to anything related to a decision problem or preference parameters. Each prior μ is just a

[13] The representation is in the form of $\min_{\mu \in \Delta(\Theta)} \ G(U(\mu), \mu)$ where $U(\mu)$ is an expected utility index for belief μ. Thus, their aggregator G takes two arguments as inputs, the conditional expected utility and the belief itself, while our aggregator V only takes in one. Therefore, strictly speaking, the uncertainty averse preference is not a special case of our set of monotone aggregators. However, it is indeed strictly increasing in its first argument, the expected utility conditional on μ. So our analysis and main result will go through for uncertainty averse preferences as well.

probability measure over Θ and is not meant to be interpreted as anything more. In this sense, this prior $\mu \in \Delta(\Theta)$ is quite different from a payoff-relevant prior belief $\pi \in \Delta(\Omega)$. The latter is a preference parameter while the former is not.

The prior-by-prior dominance condition is our desired generalization of Blackwell's garbling condition since it is the *equivalent* condition for guaranteeing a weakly higher ex-ante utility in every decision problem for every DM. Intuitively, if $\mathbf{p} \trianglerighteq_{PBP} \mathbf{p}'$, then a DM can rank them according to the Blackwell order if his preference can be represented by an aggregator that corresponds to a single subjective belief over Θ. Having such a representation is not possible in general since there is ambiguity on Θ. Our result states that $\mathbf{p}$ and $\mathbf{p}'$ can still be ranked, independent of any decision problem, when more general ambiguity preferences are considered.

Theorem 2.2. *Let $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ and $\mathbf{p}' : \Omega \times \Theta \to \Delta(S')$ be two ambiguous experiments, and let $\mathcal{V}$ be a class of aggregators such that $\mathcal{V}_{EU} \subset \mathcal{V} \subset \mathcal{V}_{Mono}$.*

The following conditions are equivalent:

1. $\mathbf{p}$ *prior-by-prior dominates* $\mathbf{p}'$.

2. *For any A, u, π and any action plan $\sigma' \in A_{S'}$, there exists $\sigma \in A_S$ such that*

$$U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) \geq U\big(\sigma', \mathbf{p}'(\cdot, \theta)\big) \text{ for all } \theta \in \Theta.$$

3. $\mathbf{p}$ *is preferred to* $\mathbf{p}'$ *in every decision problem by every decision maker whose ambiguity preference can be represented by some $V \in \mathcal{V}$. That is,* $\mathbf{p}$ *gives weakly higher ex-ante utility than* $\mathbf{p}'$ *for every A, u, π, and every $V \in \mathcal{V}$.*

The proof of Theorem 2.2 is omitted since it is a corollary of a more general theorem (Theorem B.1), formally stated in Appendix B.1 and proved in Appendix B.2, that covers the case where the auxiliary state space Θ can be any nonempty set.

Condition 1 is a statistical condition about the ambiguous experiments and independent of any decision problem or preference parameters, and condition 3 is an economical condition about the instrumental values of ambiguous experiments. Establishing the equivalence of conditions 1 and 3 helps us achieve the separation of the preferences and information structures like Blackwell's theorem.

Condition 2 states that for any action plan σ' made for $\mathbf{p}'$, there exists an action plan σ for $\mathbf{p}$ that guarantees a *weakly higher* expected utility in every auxiliary state. It is an intermediate condition for the clarification of the difference of Theorem 2.2 and Blackwell's theorem. It also highlights the importance of our use of monotone aggregators to represent ambiguity preferences since we only have weak inequality for the conditional expected utilities, unlike the equality we had in its counterpart in Blackwell's theorem (condition 2 of Theorem 2.1).

Remark. Prior-by-prior dominance is not implied by "state-by-state dominance" (which simply means that $\mathbf{p}(\cdot,\theta)$ is Blackwell more informative than $\mathbf{p}'(\cdot,\theta)$ for all $\theta \in \Theta$). To see this, consider the following example.

Example 2.4. Let $\Omega = \{\omega_1, \omega_2\}$, $S = S' = \{s_1, s_2\}$, and $\Theta = \{\theta_1, \theta_2\}$, and

$\mathbf{p}(\cdot,\theta_1) =$

	s_1	s_2
ω_1	1	0
ω_2	0	1

$\mathbf{p}(\cdot,\theta_2) =$

	s_1	s_2
ω_1	0	1
ω_2	1	0

$\mathbf{p}'(\cdot,\theta_1) =$

	s_1	s_2
ω_1	0.9	0.1
ω_2	0.1	0.9

$\mathbf{p}'(\cdot,\theta_2) =$

	s_1	s_2
ω_1	0.9	0.1
ω_2	0.1	0.9

Then $\mathbf{p}(\cdot,\theta_i)$ is strictly more informative than $\mathbf{p}'(\cdot,\theta_i)$ for $i \in \{1,2\}$, but for any belief with $\mu(\theta_1) \in (0.1, 0.9)$, $\mathbf{p}'_\mu$ is strictly more informative than $\mathbf{p}_\mu$.

Another condition that is closely related to prior-by-prior dominance is *global Blackwell dominance.*

Definition 2.9. Let $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ and $\mathbf{p}' : \Omega \times \Theta \to \Delta(S')$ be two ambiguous experiments. We say $\mathbf{p}$ *globally Blackwell dominates* $\mathbf{p}'$ if there exists a garbling $\gamma : S \to \Delta(S')$ such that $\mathbf{p}'(\cdot,\theta) = \gamma \circ \mathbf{p}(\cdot,\theta)$ for all $\theta \in \Theta$, and write $\mathbf{p} \trianglerighteq_{GB} \mathbf{p}'$.

Given the linearity of the expectation operator, $\mathbf{p} \trianglerighteq_{GB} \mathbf{p}'$ if and only there exists a single garbling $\gamma : S \to \Delta(S')$ such that

$$\mathbf{p}'_\mu = \gamma \circ \mathbf{p}_\mu \text{ for all } \mu \in \Delta(\Theta). \tag{2.10}$$

Comparing equations (2.9) and (2.10), we can see intuitively that prior-by-prior dominance is a weaker condition than global Blackwell dominance: To satisfy prior-by-prior dominance, different μ and μ' in $\Delta(\Theta)$ can correspond to different garblings γ_μ and $\gamma_{\mu'}$, but to satisfy global Blackwell dominance, one garbling γ has to work uniformly across all μ. The following proposition formalizes this intuition.

Proposition 2.3. *Let* $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ *and* $\mathbf{p}' : \Omega \times \Theta \to \Delta(S')$ *be two ambiguous experiments. If* $\mathbf{p}$ *globally Blackwell dominates* $\mathbf{p}'$*, then* $\mathbf{p}$ *prior-by-prior dominates* $\mathbf{p}'$*. The converse is not true.*

Proof. See Appendix B.2.1. □

Within our current framework, global Blackwell dominance is *not* our desired generalization of Blackwell's garbling condition since it is too strong and not necessary for always guaranteeing a weakly higher ex-ante utility. It will become more relevant when we discuss the case where u can directly depend on Θ in Section 2.5.

2.3.3 Connection and Differences with Blackwell's Theorem

Theorem 2.2 is a direct generalization of Blackwell's theorem. When Θ is a singleton, $\mathbf{p}$ and $\mathbf{p}'$ reduce to Blackwell experiments, and the prior-by-prior dominance condition reduces to the garbling condition. Prior-by-prior dominance is necessary within the small class of $\mathcal{V}_{EU}$ and it is sufficient within the large class of ambiguity preferences with $\mathcal{V}_{Mono}$. Therefore, prior-by-prior dominance is the equivalent condition corresponding to higher ex-ante utilities within the many classes of ambiguity preferences nested between them (e.g., the class of multiple prior preferences $\mathcal{V}_{MP}$ and the class of smooth preferences $\mathcal{V}_S$).

We take Blackwell's theorem as a starting point and build our result on it, but we do not replicate its proof and our result is not its trivial implication. Taking Blackwell's theorem as given, the necessity of our prior-by-prior dominance condition is relatively easy to prove: If $\mathbf{p}_\mu$ is not Blackwell more informative than $\mathbf{p}'_\mu$ for some μ, one can fix its corresponding expected utility aggregator V_μ and Blackwell's theorem guarantees the existence of some triplet (A, u, π) in which $\mathbf{p}'_\mu$ outperforms $\mathbf{p}_\mu$.

However, proving the sufficiency of our prior-by-prior dominance condition is more subtle and requires more work.[14] The main obstacle lies in the difference of condition 2 in Theorem 2.2 and its counterpart in Blackwell's theorem (condition 2 in Theorem 2.1). One quick way to prove the sufficiency of the garbling condition in Blackwell's theorem is to realize that if $p' = \gamma \circ p$, that is, if an unambiguous experiment p' is obtained by garbling p with γ, then for any action plan σ' made for p', one can obtain exactly the same expected utility under the experiment p by following the action plan $\sigma := \sigma' \circ \gamma$, since the property of compositions of stochastic operators guarantees that

$$U(\sigma, p) = U(\sigma', p'), \text{ for all } A, u \text{ and } \pi. \tag{2.11}$$

This proof method could work if we were trying to show that global Blackwell dominance is sufficient for always guaranteeing a weakly higher ex-ante utility. That is, if we assume there exists a garbling γ satisfying $\mathbf{p}_\mu = \gamma \circ \mathbf{p}'_\mu$ for all $\mu \in \Delta(\Theta)$, then for any A, u, π and action plan σ' for $\mathbf{p}'$, the action plan $\sigma := \sigma' \circ \gamma$ satisfies

$$U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) = U\big(\sigma', \mathbf{p}'(\cdot, \theta)\big), \text{ for all } \theta \in \Theta. \tag{2.12}$$

[14] This highlights another difference of our result and Blackwell's theorem in terms of their proof strategy. For Blackwell's theorem, it is easier to prove the sufficiency of the garbling condition than its necessity (for example, see the proofs of Crémer (1982) and de Oliveira (2018)). But for our result, it is the sufficiency of the prior-by-prior dominance condition that is relatively hard to prove.

This proof method does not work when we are trying to show that prior-by-prior dominance is sufficient for guaranteeing a weakly higher ex-ante utility, since the prior-by-prior dominance condition allows the garbling to depend on μ, and as argued before, it is much weaker than requiring $\mathbf{p}$ to globally Blackwell dominate $\mathbf{p}'$. Under this weaker requirement, it is generally impossible to simultaneously obtain the same conditional expected utility under $\mathbf{p}'$ and $\mathbf{p}$ in every auxiliary state θ like in equation (2.12), since there is a set of garblings and it is unclear which one should be applied to an action plan for $\mathbf{p}'$ to obtain a suitable action plan for $\mathbf{p}$.

Since the aggregator is monotone, proving the sufficiency of the prior-by-prior dominance condition does not require a condition as strong as equation (2.12). We can replace the equality in equation (2.12) with a weak inequality. This is the idea behind condition 2 in Theorem 2.2. Two key steps in proving condition 1 implies condition 2 are: (i) a careful construction of an auxiliary function capturing the differences in conditional expected utilities under $\mathbf{p}$ and $\mathbf{p}'$ and (ii) the use of a general minimax theorem to characterize the property of said auxiliary function.

2.3.4 Experiments with Independent Sources of Ambiguity

Theorem 2.2 can be applied to comparing ambiguous experiments with independent sources of ambiguity if we assume more structure to the auxiliary state space Θ. Formally, consider an auxiliary state space Θ that can be written as $\Theta_1 \times \Theta_2$ where Θ_1 and Θ_2 can be interpreted as two different aspects of the source of ambiguity. We say two ambiguous experiments $\mathbf{p} : \Omega \times \Theta_1 \times \Theta_2 \to \Delta(S)$ and $\mathbf{q} : \Omega \times \Theta_1 \times \Theta_2 \to \Delta(T)$ have independent sources of ambiguity if $\mathbf{p}$ does not depend on Θ_2 and $\mathbf{q}$ does not depend on Θ_1. More precisely, $\mathbf{p}$ and $\mathbf{q}$ have independent sources of ambiguity if there exist ambiguous experiments $\hat{\mathbf{p}} : \Omega \times \Theta_1 \to \Delta(S)$ and $\hat{\mathbf{q}} : \Omega \times \Theta_2 \to \Delta(T)$ satisfying

$$\mathbf{p}(\cdot, \theta_1, \theta_2) \equiv \hat{\mathbf{p}}(\cdot, \theta_1) \text{ for all } \theta_2 \in \Theta_2, \text{ and } \mathbf{q}(\cdot, \theta_1, \theta_2) \equiv \hat{\mathbf{q}}(\cdot, \theta_2) \text{ for all } \theta_1 \in \Theta_1.$$

We sometimes refer to $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$ as the reduced experiment of $\mathbf{p}$ and $\mathbf{q}$, respectively. With this additional structure, the prior-by-prior dominance condition relating $\mathbf{p}$ and $\mathbf{q}$ can be simplified as follows.

Lemma 2.4. *Given the above construction,* $\mathbf{p} \trianglerighteq_{PBP} \mathbf{q}$ *if and only if* $\hat{\mathbf{p}}_\mu \trianglerighteq \hat{\mathbf{q}}_\nu$ *for every pair of marginal distributions* $(\mu, \nu) \in \Delta(\Theta_1) \times \Delta(\Theta_2)$.

Proof. See Appendix B.2.2. □

That is, $\mathbf{p}$ prior-by-prior dominates $\mathbf{q}$ if and only if their reduced experiments are related by the following condition: the expected experiment $\hat{\mathbf{p}}_\mu$ is Blackwell more informative than $\hat{\mathbf{q}}_\nu$ for any $\mu \in \Delta(\Theta_1)$ and $\nu \in \Delta(\Theta_2)$. When Θ_1 and Θ_2 are identical copies of the same space, this condition becomes another variation of the prior-by-prior dominance condition: To compare the informativeness of $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$, the DM's lack of understanding of the correlation of $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$ requires him to compare all possible pairs of expected experiments $(\hat{\mathbf{p}}_\mu, \hat{\mathbf{q}}_\nu)$ corresponding to potentially different beliefs μ and ν, in contrast to the prior-by-prior dominance condition where the comparison is made between expected experiments corresponding to the same belief.

Suppose $\Theta = \Theta_1 \times \Theta_2$. Let $\mathcal{V}_{EU}$ and $\mathcal{V}_{Mono}$ denote the class of expected utility aggregators and the class of monotone aggregators over Θ, respectively. We can combine Theorem 2.2 and Lemma 2.4 to have the following result regarding the comparison of experiments with independent sources of ambiguity.

Corollary 2.5. *Let* $\mathbf{p}$, $\mathbf{q}$, $\hat{\mathbf{p}}$ *and* $\hat{\mathbf{q}}$ *be constructed as above. Suppose* $\mathcal{V}$ *is a class of aggregators such that* $\mathcal{V}_{EU} \subset \mathcal{V} \subset \mathcal{V}_{Mono}$, *then the following conditions are equivalent:*

1. $\hat{\mathbf{p}}_\mu$ *is Blackwell more informative than* $\hat{\mathbf{q}}_\nu$ *for all* $(\mu, \nu) \in \Delta(\Theta_1) \times \Delta(\Theta_2)$.

2. $\mathbf{p}$ *is preferred to* $\mathbf{q}$ *in every decision problem by every decision maker whose ambiguity preference can be represented by some* $V \in \mathcal{V}$. *That is,* $\mathbf{p}$ *gives weakly higher ex-ante utility than* $\mathbf{q}$ *for every* A, u, π, *and every* $V \in \mathcal{V}$.

In particular, the class of aggregators stated in condition 2 involves those aggregators that only aggregate over the relevant aspect of the auxiliary state space for each experiment. For example, for any subjective expected utility aggregator $V_\eta \in \mathcal{V}_{EU}$, V_η aggegrating over $\Theta_1 \times \Theta_2$ for $\mathbf{p}$ is equivalent to an expected utility aggegator over Θ_1 for $\hat{\mathbf{p}}$. Similarly, V_η aggegrating over $\Theta_1 \times \Theta_2$ for $\mathbf{q}$ is equivalent to an expected utility aggegator over Θ_2 for $\hat{\mathbf{q}}$. In this sense, condition 2 also describes the comparison of ex-ante utilities for the reduced experiments $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$.

2.4 Comparing Sets of Experiments

In this section, we study a special case where the comparison is between sets of Blackwell experiments and a DM's ex-ante utility is computed according to the maximin criterion. It is special in two ways. First, the auxiliary state space Θ is not explicitly needed at the outset for the comparison. Second, we focus on the class of ambiguity preferences corresponding to the maximin criterion. Theorem 2.2 does not apply in this special case since this class of ambiguity preferences does not nest expected utility, and we will have an informativeness order that is weaker than prior-by-prior dominance. The result is stated in Section 2.4.1, and we clarify the connection of the special case with the general comparison of ambiguous experiments in Section 2.4.2. The differences between the special case we considered and the model studied in Gensbittel, Renou, and Tomala (2015) are discussed in Section 2.4.3.

2.4.1 Result

A set of Blackwell experiments is summarized as a pair (S, P) where S is a finite set of signal realizations and P is a set of Blackwell experiments where each $p \in P$ is a mapping from Ω to $\Delta(S)$. Note that we no longer have the auxiliary state space Θ as a primitive. P should be interpreted as all unambiguous experiments that are deemed possible by the DM. We identify P as a subset of $\mathbb{R}^{|\Omega| \times |S|}$ and assume it to be closed under the Euclidean topology.[15] P can be uncountable.

Consider an individual with a finite set of actions A, a state-dependent utility function $u : \Omega \times A \to \mathbb{R}$ and a prior $\pi \in \Delta(\Omega)$. We define the individual's *ex-ante maximin utility* from a set of Blackwell experiments P to be

$$\max_{\sigma \in A_S} \min_{p \in P} U(\sigma, p), \tag{2.13}$$

where U is defined as in equation (2.1). That is, when evaluating an action plan σ for P, the DM uses the experiment in P that gives the lowest expected utility.

Let $\text{conv}(P)$ denote the convex hull of P.[16]

Theorem 2.6. *Fix two sets of experiments P and P'. The following are equivalent:*

1. *For any Blackwell experiment $p \in \text{conv}(P)$, there exists another Blackwell experiment $p' \in \text{conv}(P')$ such that p is Blackwell more informative than p'.*

2. *Every decision maker who applies Wald's maximin criterion prefers P to P' for any possible decision problem. That is, P gives weakly higher ex-ante maximin utility than P' for every A, u, and π.*

Proof. See Appendix B.2.3 □

If any of the conditions in Theorem 2.6 holds, we say P is W-more informative than P' and write $P \trianglerighteq_W P'$. To better compare $\trianglerighteq_W$ and $\trianglerighteq_{PBP}$, it is useful to define expected experiments for a set of Blackwell experiment. For any nonempty set Y, let $\Delta_0(Y)$ denote the set of probability measures over Y with finite support. Let (S, P) be a set of Blackwell experiments and $\nu \in \Delta_0(P)$. Then the *expected experiment* with respect to ν, denoted by $P_\nu : \Omega \to \Delta(S)$, is defined by $P_\nu := \sum_{p \in P} p\nu(p)$. With this notion, condition 1 can be rephrased as

[15] P is a bounded set since for every element $p \in P$, every entry of p is bounded in $[0, 1]$. Hence its closedness implies its compactness.

[16] The convex combination is taken component-wise, that is, for any $p, q \in P$, $tp + (1-t)q : \Omega \to \Delta(S)$ is defined by $(tp + (1-t)q)(s \mid \omega) = tp(s \mid \omega) + (1-t)q(s \mid \omega)$.

1'. For all $\nu \in \Delta_0(P)$, there exists some $\lambda \in \Delta_0(P')$ such that the expected experiment P_ν is Blackwell more informative than P'_λ.

Comparing condition 1' and the prior-by-prior dominance condition, we can see intuitively that the W-more informative condition is somewhat less restrictive than the prior-by-prior dominance condition as the latter requires the Blackwell order to hold for every pair of expected experiments corresponding to the same $\mu \in \Delta(\Theta)$. We obtain this less restrictive condition mainly because we have a more restrictive class of preferences. DMs who apply the maximin criterion essentially go through all possible priors over Θ and then focus only to the worst possible ones, and thus ignoring the effect of other priors.

2.4.2 Connection with Ambiguous Experiments

Ambiguous experiments and sets of Blackwell experiments are different primitives. However, there is a natural correspondence between them and they are essentially equivalent when we focus our attention to DMs who use Wald's maximin criterion.

To see the correspondence, suppose we have two sets of Blackwell experiments (S, P) and (S', P'), then we can construct an auxiliary state space Θ, and two ambiguous experiments $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ and $\mathbf{p}' : \Omega \times \Theta \to \Delta(S')$ corresponding to (S, P) and (S', P'), respectively, in the following way:

$$\begin{aligned} \Theta &:= P \times P' \\ \mathbf{p}(s \mid \omega, p, p') &:= p(s \mid \omega),\ \forall (s, \omega, p, p') \in S \times \Omega \times P \times P';\ \text{and} \\ \mathbf{p}'(s' \mid \omega, p, p') &:= p'(s' \mid \omega),\ \forall (s', \omega, p, p') \in S' \times \Omega \times P \times P' \end{aligned} \tag{2.14}$$

That is, the situation with two sets of Blackwell experiments can be re-interpreted as following: the DM wants to compare P and P' but has no additional knowledge on how their elements are correlated. He therefore creates $\Theta = P \times P'$ as the auxiliary state space, as if he is willing to entertain each and every pair of the elements $(p, p') \in P \times P'$ to be the actual Blackwell experiments he will be facing.

Now suppose we have two ambiguous experiments $\mathbf{p}$ and $\mathbf{p}'$, then we can construct two sets of Blackwell experiments

$$P := \{\mathbf{p}(\cdot, \theta) \mid \theta \in \Theta\} \quad P' := \{\mathbf{p}'(\cdot, \theta) \mid \theta \in \Theta\} \tag{2.15}$$

That is, each set of Blackwell experiments is created by collecting the Blackwell experiments in each auxiliary state in its corresponding ambiguous experiment.

Given these constructions in (2.14) and (2.15), comparisons of sets of Blackwell experiments using maximin criterion are equivalent to comparisons of ambiguous experiments for DMs with the Wald aggregator V_W.[17]

Proposition 2.7. *If* $\mathbf{p}$, $\mathbf{p}'$, P *and* P' *satisfies one of the constructions given by* (2.14) *or* (2.15) *with* P *and* P' *being finite,*[18] *then for any* A, u, π, *and any action plans* σ *and* σ',

$$V_W\left(\left\{U\left(\sigma, \mathbf{p}(\cdot, \theta)\right)\right\}_{\theta \in \Theta}\right) = \min_{p \in P} U(\sigma, p), \qquad V_W\left(\left\{U\left(\sigma', \mathbf{p}'(\cdot, \theta)\right)\right\}_{\theta \in \Theta}\right) = \min_{p' \in P'} U(\sigma', p').$$

Proof. See Appendix B.2.4. □

In general, comparisons of sets of Blackwell experiments is a special case of comparisons of ambiguous experiments, and the resulting informativeness order induced by the W-more informative condition is different with that induced by the prior-by-prior dominance condition. This is because the comparisons of sets of Blackwell experiments are only possible for the specific class of decision makers who uses the maximin criterion. This class of decision makers corresponds to the Wald aggregator V_W, but $\{V_W\}$ is not a superset of $\mathcal{V}_{EU}$, that is, the class of expected utility aggregators is not a subclass of the Wald aggregator. Our argument to get the necessity of

[17] Recall that $V_W : \mathbb{R}^\Theta \to \mathbb{R}$ is defined by $V_W(f) = \min_{\theta \in \Theta} f(\theta)$.

[18] The finiteness assumption in Proposition 2.7 is only for ease of exposition, a more general proposition (Proposition 2.9) without this assumption is stated and proved in the Appendix.

prior-by-prior dominance in the general model is through its necessity for the class of expected utility aggregators. Now that $\{V_W\}$ no longer includes $\mathcal{V}_{EU}$, the prior-by-prior dominance is no longer necessary. The following proposition states that prior-by-prior dominance is a stronger requirement than being W-more informative.

Proposition 2.8. *Let P and P' be two finite sets of Blackwell experiments. Define $\mathbf{p}$ and $\mathbf{p}'$ as in* (2.14)*, then $\mathbf{p} \trianglerighteq_{PBP} \mathbf{p}'$ implies $P \trianglerighteq_W P'$. The converse is not true.*

Proof. Suppose it is not the case that $P \trianglerighteq_W P'$, then there exists (A, u, π) in which P' yields strictly higher ex-ante maximin utility than P. Combine this triplet (A, u, π) with the aggregator V_W, we find a quadruple (A, u, π, V_W) in which $\mathbf{p}'$ yields strictly higher ex-ante utility than $\mathbf{p}$. Thus, by Theorem 2.2, $\mathbf{p}$ does not prior-by-prior dominates $\mathbf{p}'$. Taking the contrapositive completes the proof.

To see the converse is not true, consider the following example:

$\Omega = \{\omega_1, \omega_2\}$, $S = S' = \{s_1, s_2\}$, and

$$P = \left\{ \begin{array}{|c|c|c|} \hline & s_1 & s_2 \\ \hline \omega_1 & 0.9 & 0.1 \\ \hline \omega_2 & 0.1 & 0.9 \\ \hline \end{array} \right\}, \; P' = \left\{ \begin{array}{|c|c|c|} \hline & s_1 & s_2 \\ \hline \omega_1 & 1 & 0 \\ \hline \omega_2 & 0 & 1 \\ \hline \end{array}, \begin{array}{|c|c|c|} \hline & s_1 & s_2 \\ \hline \omega_1 & 0 & 1 \\ \hline \omega_2 & 1 & 0 \\ \hline \end{array} \right\}.$$

Then $P \trianglerighteq_W P'$, but no matter how we defined the auxiliary state space Θ, we do not get prior-by-prior dominance. □

Modeling ambiguous experiments as sets of unambiguous experiments has its advantage: The framework for Blackwell's theorem can be applied without much modification. But such an advantage comes with costs tightly connected with the maximin criterion: DMs' potential subjective belief over possible unambiguous experiments are completely ignored due to the extreme ambiguity attitude. Although the result for comparisons of sets of experiments is interesting in itself, we may naturally also be interested in environments that can allow for richer ambiguity preferences.

Figure 2.1 summarizes the results we have so far.

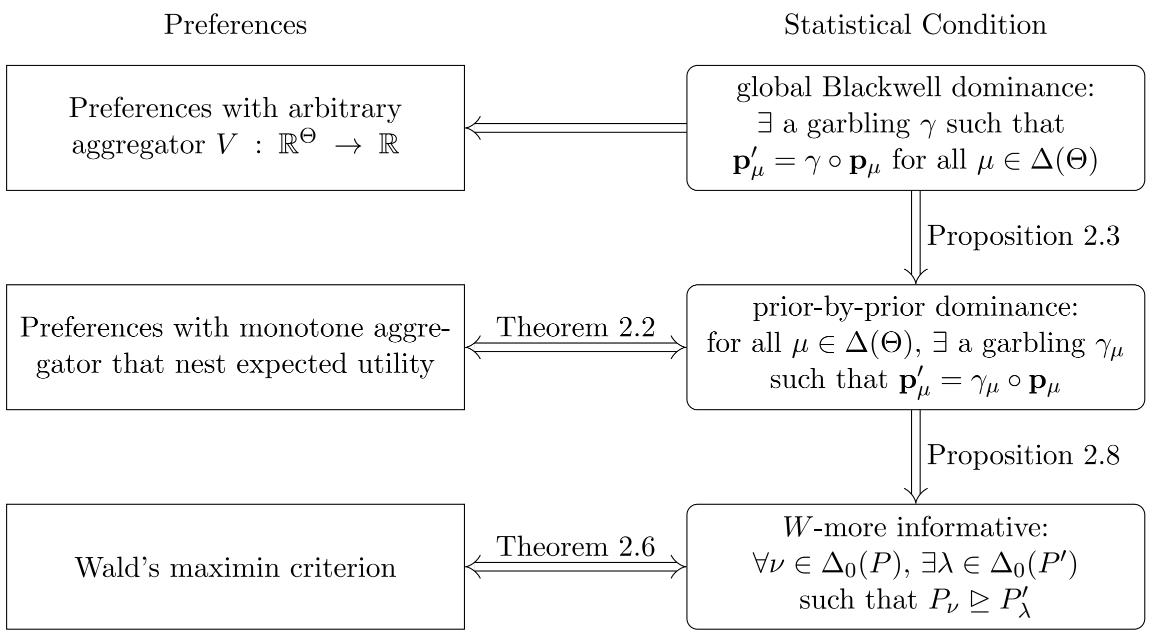

Figure 2.1: Summary of informativeness orders.

2.4.3 Connection with the Literature

Our Theorem 2.6 is closely related to an existing result in Gensbittel, Renou, and Tomala (2015) (henceforth GRT). Although arriving at seemingly similar results, there are subtle but important differences between our model and result for comparisons of sets of experiments (W-more informative) and the model and result presented in GRT due to a major methodological difference: We model Blackwell experiments as collections of conditional distributions $p : \Omega \to \Delta(S)$, while GRT model them as joint distributions $k \in \Delta(S \times \Omega)$ over the product space of the payoff-relevant states and the signal realizations. These two approaches are equivalent in the framework of Blackwell's theorem. However, these two approaches are no longer equivalent when ambiguous experiments are studied.

GRT model ambiguous experiments as *compact* and *convex* subsets of $\Delta(S \times \Omega)$. The benefit of this modeling approach is that one does not need to specify the source of ambiguity. That is, this method can deal with two special cases: An unambiguous prior combined with a set of Blackwell experiments (in our sense, i.e., collections of conditional distributions), or a set of priors combined with a single Blackwell experiment. To get this benefit, however, the inevitable cost is that the DM's prior over Ω is no longer a part of the DM's preference parameters, but a part of the experiment. For example, we can fix a Blackwell experiment $p : \Omega \to \Delta(S)$ and pair them with two different (compact and convex) sets of priors, $\Pi_1 \subsetneq \Pi_2$ on Ω. Then we can find examples of Π_1 and Π_2 such that the set of joint distributions induced from a Blackwell experiment p and Π_1 is deemed "more informative than" the set of joint distributions induced from p and Π_2 by GRT's criterion. For a simple example, consider a decision problem with a single action $A = \{a\}$. Then a larger set of priors necessarily result in lower payoff as the minimum is taken over a larger set. Thus, their criterion compares the informativeness and ambiguity associated with the experiment, but also the ambiguity of the specified set of priors on Ω.

Therefore, our criteria differ, and both are useful in different circumstances: Our criterion is more useful when the modeler cannot observe the decision maker's prior over Ω, while GRT's criterion is more suitable when those priors can be observed.

2.5 Discussion and Extensions

In this section, we consider relaxing two of the assumptions we previously made. The first is that u does not depend on Θ, that is, the auxiliary state does not directly affect the payoff of the DM. And the second is that there is no prior ambiguity, that is, the decision problems we consider all involve unambiguous prior over Ω. Table 2.1 is a roadmap for our exercises on relaxing these two assumptions.

2.5.1 Direct Dependence of Payoffs on Θ

In this section, we consider individuals whose payoff may directly depend on the auxiliary state space Θ. That is, the individual's state dependent utility function

Table 2.1: Roadmap of Extensions

	Section 2.3	2.5.1	2.5.2
Dependence of u on Θ	No	Yes	No
Prior ambiguity on Ω	No	Yes/No	Yes
Informativeness order	prior-by-prior dominance	global Blackwell	unclear

is $u : \Omega \times \Theta \times A \to \mathbb{R}$. With this change in setup, an ambiguous experiment $\mathbf{p}$ is preferred to $\mathbf{p}'$ by every DM for every decision problem if and only if $\mathbf{p}$ globally Blackwell dominates $\mathbf{p}'$, which is equivalent to assuming that $\mathbf{p}$ is Blackwell more informative than $\mathbf{p}'$ viewing $\Omega \times \Theta$ as the payoff-relevant state space.

To see why this is the case, note that once the DM's payoff u depends directly on Θ, we can expand the payoff-relevant state space from Ω to $\Omega \times \Theta$. With this expansion, the ambiguity previously in the information structures is transformed into the ambiguity in the priors over Θ. That is, the decision maker faces a Blackwell experiment although there is ambiguity in his payoff-relevant prior. This problem is the center of the study in Li and Zhou (2016), where they show that Blackwell order continues to be valid when the decision makers possess uncertainty averse preferences to deal with the ambiguity in the payoff relevant space.[19] The intuition behind this result is as follows. If one experiment $\mathbf{p}$ globally Blackwell dominates another $\mathbf{p}'$, then any probability distribution over the actions *conditional* on the states ω and θ, $\lambda(a \mid \omega, \theta)$, that can be induced from some action plan for $\mathbf{p}'$ can be replicated under $\mathbf{p}$ by applying the garbling that transforms $\mathbf{p}$ to $\mathbf{p}'$, and this replication does not depend on the prior over the payoff relevant state space. Therefore, as long as the payoffs $u(\omega, \theta, a)$ are aggregated through some aggregator that respects monotonicity (e.g., the aggregator for the uncertainty averse preferences), the Blackwell informativeness order will be valid.

[19] See Theorem 1 of Li and Zhou (2016) for more details.

2.5.2 Prior Ambiguity

In this section, we consider decision scenarios in which prior ambiguity (i.e., ambiguity over Ω) is present. As the case where u depends on Θ and prior ambiguity over Ω are both present has already been covered in Section 2.5.1, we focus on the case where u does not depend directly on Θ in this section.

Formally, we consider an individual with a finite set of actions A, a state-dependent utility function $u : \Omega \times A \to \mathbb{R}$ and a monotone aggregator $\hat{V} : \mathbb{R}^{\Omega\times\Theta} \to \mathbb{R}$ that captures his ambiguity attitude. Notice that $\hat{V}$ is more general than the aggregators V we considered in Section 2.3 since $\hat{V}$ is aggregating over functions in $\mathbb{R}^{\Omega\times\Theta}$ while V is only aggregating over functions in $\mathbb{R}^{\Theta}$.

For such an individual, the expected utility conditional on θ and ω for ambiguous experiment $\mathbf{p}$ and action plan σ is

$$\hat{U}(\sigma, \mathbf{p}(\cdot, \theta), \omega) := \sum_{s\in S}\sum_{a\in A} \mathbf{p}(s \mid \omega, \theta)\sigma(a \mid s)u(\omega, a) \tag{2.16}$$

The *ex-ante utility* for the individual is

$$\max_{\sigma\in A_S} \hat{V}\left(\left\{\hat{U}\big(\sigma, \mathbf{p}(\cdot, \theta), \omega\big)\right\}_{(\theta,\omega)\in\Theta\times\Omega}\right) \tag{2.17}$$

With this new formulation involving more general aggregators $\hat{V}$, the payoffs $u(\omega, a)$ can be aggregated in a more general way comparing to equation (2.3). For example, it can nest the case where the DM has a set of priors over Ω, first aggregating each $u(\omega, a)$ according to the multiple prior preference conditional on each θ and then aggregating the conditional utilities into the ex-ante utility.

The following proposition illustrates the results when prior ambiguity is present.

Proposition 2.9. *Let* $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ *and* $\mathbf{p}' : \Omega \times \Theta \to \Delta(S')$ *be two ambiguous experiments. Consider the following 4 conditions:*

1. $\mathbf{p}$ *globally Blackwell dominates* $\mathbf{p}'$*, that is,* $\mathbf{p} \trianglerighteq_{GB} \mathbf{p}'$*.*

2. For any A, u and any $\sigma' \in A_{S'}$, there exists $\sigma \in A_S$ such that

$$\hat{U}(\sigma, \mathbf{p}(\cdot, \theta), \omega) \geq \hat{U}(\sigma', \mathbf{p}'(\cdot, \theta), \omega), \ \forall (\theta, \omega) \in \Theta \times \Omega.$$

3. $\mathbf{p}$ gives a weakly higher ex-ante utility than $\mathbf{p}'$ for every A, u and $\hat{V}$.

4. $\mathbf{p}$ prior-by-prior dominates $\mathbf{p}'$, that is, $\mathbf{p} \trianglerighteq_{PBP} \mathbf{p}'$.

Then, $1 \implies 2 \iff 3 \implies 4$.

Proof. See Appendix B.2.5. □

That is, in terms of guaranteeing higher ex-ante utility for every decision maker (condition 3), global Blackwell dominance is sufficient and prior-by-prior dominance is necessary. Although condition 2 is equivalent to condition 3, it is less ideal than what we desire as an equivalence condition because it is a somewhat high-order condition imposed on all possible A and u.

The following example illustrates that when prior ambiguity is present, global Blackwell dominance is not necessary for guaranteeing higher ex-ante utility. That is, condition 3 does not imply condition 1.

Example 2.5. Consider $\Omega = \{\omega_1, \omega_2\}$, $\Theta = \{\theta_1, \theta_2\}$, $S = S' = \{s_1, s_2\}$, with $\mathbf{p}$ and $\mathbf{p}'$ defined as follows:

$\mathbf{p}(\cdot, \theta_1) =$

	s_1	s_2
ω_1	1	0
ω_2	0	1

$\mathbf{p}(\cdot, \theta_2) =$

	s_1	s_2
ω_1	1	0
ω_2	0	1

$\mathbf{p}'(\cdot, \theta_1) =$

	s_1	s_2
ω_1	1	0
ω_2	0	1

$\mathbf{p}'(\cdot, \theta_2) =$

	s_1	s_2
ω_1	0	1
ω_2	1	0

That is, $\mathbf{p}$ is the unambiguous fully revealing experiment while $\mathbf{p}'$ is fully revealing in both θ_1 and θ_2 but sends opposite signals in different auxiliary states. It is clear that $\mathbf{p}$ prior-by-prior dominates $\mathbf{p}'$ but $\mathbf{p}$ does not globally Blackwell dominates $\mathbf{p}'$.

Proposition 2.10. *As constructed in Example 2.5,* $\mathbf{p}$ *gives weakly higher ex-ante utility than* $\mathbf{p}'$ *for every* A*,* u*, and* $\hat{V}$*.*

Proof. See Appendix B.2.6. □

2.6 Conclusion

In this chapter, we study comparisons of ambiguous experiments and establish informativeness orders over ambiguous experiments as generalizations of Blackwell's theorem. For general ambiguous experiments modeled as mappings from an auxiliary state space to the space of unambiguous experiments, the informativeness order is induced by the prior-by-prior dominance condition. One ambiguous experiment $\mathbf{p}$ prior-by-prior dominates another $\mathbf{p}'$ if their expected experiments satisfy $\mathbf{p}_\mu$ being Blackwell more informative than $\mathbf{p}'_\mu$ for every possible belief μ over the auxiliary state space. This informativeness order is robust among any monotone ambiguity preference that nests expected utility. For the special case of comparing sets of unambiguous experiments evaluated by the maxmin criterion, the informativeness order is characterized by being Wald-more informative. One set of unambiguous experiment P is Wald-more informative than another P' if for every unambiguous experiment in the convex hull of P, there exists an unambiguous experiment in the convex hull of P' that is Blackwell less informative.

An interesting question that remains open is whether or not prior-by-prior dominance is sufficient for guaranteeing higher ex-ante utility when ambiguity on Ω is present. Another potentially fruitful avenue for future research is to study the impact of relaxing the dynamic consistency assumption within the class of ambiguity preferences more general than the maxmin expected utility.

Chapter 3

Persuasion with a Constrained Signal Space

3.1 Introduction

There has been an explosion of papers on Bayesian persuasion since Kamenica and Gentzkow (2011, henceforth KG) coined the term. Most of the papers in this line of literature adopts (mostly implicitly) the assumption that the Sender can choose signals that have *arbitrary* joint distributions with the underlying state of the world, i.e., there is no *ex ante* constraint on the signals/information structure he can choose. This chapter is motivated by the observation that this assumption is a bit too strong to maintain in many applications. For example, in the leading example of KG, a prosecutor (Sender) tries to persuade a judge (Receiver) to convict the defendant by conducting an investigation to change the judge's belief about how likely the defendant is guilty. In the optimal persuasion mechanism, the prosecutor would want an investigation that generates a deterministic signal for a guilty defendant and a stochastic signal for an innocent one. In particular, the fully revealing signal is assumed to be feasible to the prosecutor. In many applications, a fully revealing signal may be *impossible* to obtain due to technological constraints. To illustrate this point, consider the same prosecutor-judge example. It might be impossible for the prosecutor to have an investigation that completely tells apart a guilty defendant from an innocent defendant due to the following reasons: (i) The precision of forensic tests may subject to technological constraints; (ii) The prosecutor may not be capable of asking the most relevant questions to witnesses; (iii) Even the best expert witness may not have definitive answers to certain questions. When one or more of the above happens, the fully revealing signal is not feasible for the prosecutor. In that sense, the prosecutor's problem in reality is a problem of persuasion with a constrained signal space.

It is not obvious how one should model such constraints on the signal space. Different approaches have been employed in several recent papers. The first one can be summarized as a "noisy channel" approach (e.g., Tsakas and Tsakas (2018), Le Treust and Tomala (2019)). In this approach, the Sender can only communicate with the Receiver through a *noisy* channel, in which there is a fixed probability an intended input message of the Sender results in a different message being sent to the Receiver. The second one can be summarized as an "entropy reduction" approach. In this approach, there is an exogenous upper bound on the reduction in expected entropy achieved by the Sender's signal. Le Treust and Tomala (2019) provides a nice comparison of the persuasion outcomes of the first two approaches. Ichihashi (2019) employs a third approach. In their model, there is an exogenous signal serving as the upper-bound of the informativeness (in the sense of Blackwell (1953)) of the signals the Sender can use. It is closely related to the first approach but has some subtle differences. The approach we take in this chapter is an "α-constraint." Roughly speaking, the α-constraint dictates that no signal realization can be very informative either about a state being the true state, or about a state not being the true state. More precisely, the probability of any signal realization being sent out conditional on any state of the world must sit in $[\alpha, 1-\alpha]$, instead of $[0, 1]$. All four approaches described above are *ex ante* different, in the sense that they exclude different signals from the same signal space in general.[1]

With the appropriate constraint in place, one would naturally ask the following questions: what is the optimal persuasion mechanism/signal/information structure for a constrained Sender? What is the corresponding persuasion payoff? How much worse is it for the Sender comparing to the standard persuasion model?

3.1.1 Preview of Results

To answer these questions, one first needs a method to "solve" the constrained persuasion problem. My first main result simplifies the analysis by showing that it is without loss of generality to consider the signals that have a signal realization space no bigger than Receiver's action space. In other words, the "action-recommendation" approach is still feasible in terms of solving this constrained problem. However, it is different from previous results in the literature (e.g., KG for the standard single-sender-single-receiver case, and Bergemann and Morris (2016) for a more general single-sender-multiple-receiver setting), which states that having a message space (interchangeably, a signal realization space) exactly equal to Receiver's action space can always be optimal. In our model, the Sender can *strictly* benefit from reducing the size of his message space, since this helps relaxing the constraint and gives him the ability to make his messages more informative. Therefore, having a message as a proper subset of the Receiver's action space can be uniquely optimal, i.e., it performs better than having the whole action space as the message space as suggested in previous studies.

With this result, we give an algorithm to solve for general constrained persuasion problems that has finite state space and action space. The algorithm goes as follows: fix a subset of Receiver's actions space that has two or more elements and use it as the set of actions Sender wants to recommend. Solve the standard linear programming

[1] We will have a formal and more detailed discussion of these different approaches in Section 3.5.

problem involving this set of action recommendations and record its corresponding persuasion payoff. Repeat until exhausting all subsets of Receiver's action space, and the one with the highest persuasion payoff is our desired solution. Although every sub-problem is a linear programming problem and easy to solve, the whole problem is not, since the number of sub-problems grows exponentially fast. The computational issue is not severe when the action space is not too large.

3.1.2 Related Literarture

This chapter aims to contribute to the literature on single-sender-single-receiver Bayesian persuasion problems by imposing an exogenous constraint on signals the Sender can choose.

As mentioned above, several recent papers are close to ours but employ different approaches on modeling the constraint. Tsakas and Tsakas (2018) uses a noisy channel approach and focuses on the Receiver's preference over the noisiness of the channel and provide an example where the Receiver is strictly better off under a more noisy communication channel, which roughly corresponds to the case where the Sender is more constrained. Le Treust and Tomala (2019) gives an information-theoretic perspective on the problem by consider noisy channels and multiple copies of persuasion problems. They consider both the noisy channel approach and the entropy reduction approach. They provide a powerful tool based on convex optimization to solve for the belief-splitting problem with the entropy-reduction constraint and proves that the optimal persuasion payoff under the entropy reduction approach is a tight bound on the optimal persuasion payoff under the noisy channel approach. Doval and Skreta (2018a) builds on Le Treust and Tomala (2019) and has an ambitious goal to provide a tool to solve persuasion problems with general constraints on the posterior beliefs (entropy reduction is one specific constraint). Ichihashi (2019) characterize the lower and upper bound on feasible persuasion payoffs for the Sender when his choice of signal must be less informative than some exogenously given signal, by cleverly using the dual approach where the role of the Sender and the Receiver are switched. Our work here is different with the papers above in considering a different constraint on the signal space.

Another relevant line of literature is persuasion with rational inattentive receivers, e.g., Bloedel and Segal (2018), Lipnowski, Mathevet, and Wei (2019), Wei (2018). In these papers, the Sender is not constrained, but processing information is costly for the Receiver. Therefore, after the Sender commits to a signal structure but before any signal realization is sent, the Receiver chooses an *attention strategy* which specifies how "precise" she wants to learn about the signal realization. This modification makes the mathematical problem similar to what we have here, but the motivation is different. They consider a story in which the Receiver may not want to pay full attention since it is too costly, while we consider a story in which the Sender's hands are tied so his signal cannot be very informative in the first place.

Many papers regarding persuasion games consider *endogenous* constraint on signal spaces, in the sense that these constraints come from optimality conditions in the model (persuasion/information design is usually just one component in these models). For example, Galperti and Perego (2019) considers persuasion with a network of receivers. And the constraint on the Sender's signal space comes from the information flow in the network. Dworczak (2017) and Doval and Skreta (2018b) are examples where a single model incorporates both a mechanism design problem and an information design problem, and the incentive compatibility constraints from the mechanism design problem impose extra constraints on the Sender's signal space. In this chapter, we consider a simpler problem where the constraint on the Sender's signal space is exogenously given.

The rest of the chapter is organized as follows: Section 3.2 revisits the prosecutor judge example and provides a complete characterization of its solution under the

α-constraint. Section 3.3 lays out the general model and presents the main result. Section 3.4 provides a characterization of feasible distribution of posteriors for a special case. Section 3.5 briefly discusses other existing approaches on modeling constraints on information structure. Section 3.6 concludes and discusses some possible next steps.

3.2 An Example

We continue to use the leading example from KG.

A prosecutor (Sender, she) is trying to persuade a judge (Receiver, he) to convict a defendant. The judge *must* choose between acquitting and convicting the defendant. There are two states of the world, the defendant is either guilty or innocent. The judge prefers to be just. He gets utility 1 for convicting the guilty and acquitting the innocent, and he gets utility 0 otherwise. The prosecutor prefers conviction. She gets utility 1 if the judge convicts and utility 0 if the judge acquits, whether the defendant is guilty or innocent. Both parties share a correctly specified prior belief $\mathsf{P}(\text{guilty}) = \mu < 0.5$, so as an expected-utility-maximizer, the judge will choose to acquit the defendant with this prior belief.

However, before the judge makes a decision, the prosecutor can conduct an investigation and is required by law to report its full outcome.[2] Formally, an investigation consists of a set, M, collecting all possible outcomes of the investigation and distributions $\pi(\cdot \mid \text{guilty})$ and $\pi(\cdot \mid \text{innocent})$ over M. The prosecutor chooses M and π and must honestly report the realized outcome to the judge. The key assumption is that the prosecutor cannot distort or conceal information once the investigation is completed and an outcome is produced.

[2] As argued in KG, one can think of the choice of the investigation as consisting of the decisions on whom to subpoena, what forensic tests to conduct, what questions to ask an expert witness, etc.

3.2.1 Introducing the α-constraint

To formally define the α-constraint, suppose the prosecutor has chosen some finite set M as the set of signal realizations (we will call this the *message space* henceforth), say, $M = \{m_1, \ldots, m_n\}$. In KG, the distributions $\pi(\cdot \mid \text{guilty})$ and $\pi(\cdot \mid \text{innocent})$ are feasible to the prosecutor as long as they are some legitimate probability distribution over the message space M, i.e.,

$$\pi(\cdot \mid \text{guilty}), \pi(\cdot \mid \text{innocent}) \in \Delta(M) := \left\{ (p_1, \ldots, p_n) \mid \sum_{i=1}^{n} p_i = 1,\ p_i \in [0,1],\ \forall i \right\}.$$

A fully revealing signal in this case could be defined by $M = \{g, i\}$ and

$$\pi(g \mid \text{guilty}) = 1,\ \pi(i \mid \text{guilty}) = 0,\ \pi(g \mid \text{innocent}) = 0,\ \pi(i \mid \text{innocent}) = 1.$$

As argued in the introduction, such a fully revealing signal may not be feasible to the prosecutor due to various technological constraints. To model this, we introduce the α-constraint as following:

Definition 3.1 (Constrained Signal Space)**.** In the prosecutor-judge example, once the message space M is chosen, the prosecutor can only choose her signals $\pi(\cdot \mid \text{guilty})$ and $\pi(\cdot \mid \text{innocent})$ from the following constrained space:

$$\Delta_\alpha(M) := \left\{ (p_1, \ldots, p_n) \mid \sum_{i=1}^{n} p_i = 1,\ p_i \in [\alpha, 1-\alpha] \right\},\ \text{with } \alpha \in \left(0, \frac{1}{n+1}\right].$$ [3]

In general, suppose we have a finite state space Θ and a message space M, the Sender can only choose signal π that satisfies:

$$\pi(m \mid \theta) \in [\alpha, 1-\alpha],\ \forall m \in M \text{ and } \forall \theta \in \Theta.$$

[3] This restriction is for ease of exposition. In particular, it guarantees that a message space $M = A$ is feasible, that is, it is possible for the Sender to use messages as action recommendations.

That is, the probability that any message m is sent out when the true state is θ can neither be too close to 0 nor be too close to 1. The assumption that α does not depend on θ is for the ease of exposition, the main result (Proposition 3.1) still holds even if we allow the constraint α to depend on the state θ.

This constraint is not as artificial as it seems at first glance. It can be micro-founded through a story of noisy communication channel, as in Le Treust and Tomala (2019) and Tsakas and Tsakas (2018), or through a story of the existence of some pre-determined upper bound of informativeness, as in Ichihashi (2019). Detailed discussions are postponed to Section 3.5.

3.2.2 Solving the Example

As we will show later for the general model, it is without loss of generality to consider the message space M with $|M| \leq |A|$. In this example, the action space is binary, i.e., $|A| = 2$. But having $|M| = 1$ means there is only one message can be sent, which corresponds to a completely uninformative signal and is not optimal for the Sender (not to mention it violates the α-constraint). Therefore, to get the optimal information structure, we need only to consider message space M with $|M| = |A| = 2$. Let $M = \{g, i\}$. One can interpret message g as a recommendation that the defendant is guilty and the judge should convict and message i as a recommendation that the defendant is innocent and the judge should acquit.

Now the prosecutor-judge example with a constrained prosecutor can be summarized in the following maximization problem:

$$\max_{\{\pi(g|\text{guilty}),\pi(g|\text{innocent})\}} \quad \mu \times \pi(g \mid \text{guilty}) + (1-\mu) \times \pi(g \mid \text{innocent})$$

$$\text{s.t.} \quad \mathsf{P}(\text{guilty} \mid g) = \frac{\mu \times \pi(g \mid \text{guilty})}{\mu \times \pi(g \mid \text{guilty}) + (1-\mu) \times \pi(g \mid \text{innocent})} \geq \frac{1}{2} \quad \text{(obedience to } g\text{)}$$

$$\mathsf{P}(\text{innocent} \mid i) = \frac{(1-\mu) \times (1-\pi(g \mid \text{innocent}))}{\mu \times (1-\pi(g \mid \text{guilty})) + (1-\mu) \times (1-\pi(g \mid \text{innocent}))} \geq \frac{1}{2} \quad \text{(obedience to } i\text{)}$$

$$\pi(g \mid \text{guilty}), \pi(g \mid \text{innocent}) \in [\alpha, 1-\alpha] \quad \text{(feasibility)}$$

That is, the prosecutor is maximizing the probability that message g is sent subject to the *obedience constraints* that it is indeed optimal for the judge to follow the action recommendations, and the *feasibility constraints* that the signal satisfies the α-constraint.

This is a standard linear programming problem. The solution is summarized in Figure 3.1 below.

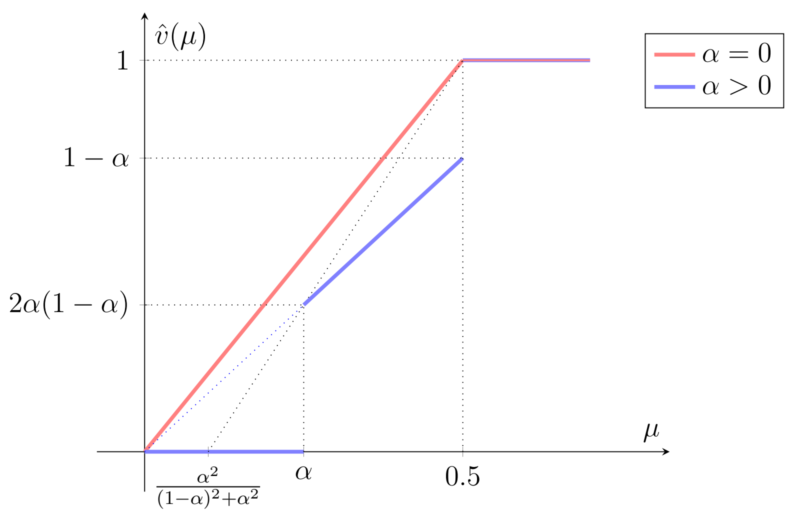

Figure 3.1: Persuasion payoff $\hat{v}(\mu)$ as a function of the prior belief μ.

In Figure 3.1, the red kinked line represents the optimal persuasion payoff for an unconstrained prosecutor, which is given by directly applying the concavification result from KG and Aumann and Maschler (1995).

The optimal persuasion payoff for a constrained prosecutor is denoted by the blue line segments. Note that it has two regions:

1. If $\mu < \alpha$, that is, the common prior is so extreme that every party believes that it is very unlikely the defendant is guilty, then due to the constraint on the accessible signals, there is simply no message (neither g nor i) that can make the judge's posterior belief reach a half. Hence in this region, the prosecutor's ex ante expected payoff stays at zero and is not improved at all.

2. If $\mu \geq \alpha$, then it is possible to adjust the judge's belief to a half, and the optimal signal structure is

$$\pi(g \mid \text{guilty}) = 1 - \alpha \quad \text{and} \quad \pi(g \mid \text{innocent}) = \frac{\mu}{1-\mu}(1-\alpha),$$

   and the ex ante expected payoff of the prosecutor is $2(1-\alpha)\mu$. That is, a message g would raise the prosecutor's posterior to a half, but the persuasion is not as efficient as when the signal is not constrained (in the sense that the α-constraint lowers the total probability of message g being sent).

This pattern (i.e., the existence of these two different regions) will persist even in more general persuasion problems. Now we are ready for the general model.

3.3 General Model

There are two players, a Sender S (she) and a receiver R (he). A constrained persuasion problem is defined as

$$\mathcal{P} := (\Theta, \mu, A, u_S, u_R, \alpha),$$

where Θ is the state space, $\mu \in \Delta\Theta$ is the correctly specified prior belief, A is the action space of the receiver, $u_S : A \times \Theta \to \mathbb{R}$ is a continuous function representing

Sender's utility, $u_R : A \times \Theta \to \mathbb{R}$ is a continuous function representing receiver's utility, and α represents the constraint on the information structure/signals that are available to the Sender. For convenience, we make some regularity assumptions.

Assumption 3.1 (Regularity)**.**

- Suppose Θ and A are finite, with $|\Theta| \geq 2$ and $|A| \geq 2$.
- Suppose the prior belief μ has full support, that is, $\mu(\theta) > 0$ for all $\theta \in \Theta$.
- Suppose for all $a \in A$, there exists some $\mu \in \Delta\Theta$ that a is an optimal action for R under belief μ. That is, for any $a \in A$, there exists a belief $\mu \in \Delta\Theta$ such that $\sum_{\theta \in \Theta} \mu(\theta) u_R(a, \theta) \geq \sum_{\theta \in \Theta} \mu(\theta) u_R(a', \theta)$ for all $a' \in A$.
- Suppose $\alpha \leq \frac{1}{|A|+1}$.

The motivation for Assumption 3.1 is straightforward. Assuming finite Θ and A is to avoid the potential complexity associated with probability distributions over an infinite space. Assuming $|\Theta|, |A| \geq 2$ guarantees the problem is not degenerate. μ having full support means there is no vacuous state, and every $a \in A$ is optimal for some belief guarantees that there is no vacuous action.

For a fixed constrained persuasion problem $\mathcal{P}$, the Sender chooses an information structure, i.e., a pair (M, π) where M is a finite message space and $\pi : \Theta \to \Delta M$ specifies the probability distribution over M condition on each state. Once the message space M is chosen, the α-constraint dictates that

$$\pi(m \mid \theta) \in [\alpha, 1 - \alpha], \ \forall m \in M \text{ and } \forall \theta \in \Theta.$$

After the Sender chooses and commits to an information structure (M, π), a message m is realized according to the conditional distribution $\pi(\cdot \mid \theta)$. R observes this message without observing the true state θ and updates his belief using Bayes' rule.

Let $\mu_m \in \Delta\Theta$ denote the posterior after observing message m, then by Bayes' rule:

$$\mu_m(\theta) = \frac{\mu(\theta)\pi(m \mid \theta)}{\sum_{\theta' \in \Theta} \mu(\theta')\pi(m \mid \theta')}, \ \forall \theta \in \Theta.$$

Based on this posterior belief μ_m, the receiver chooses an action that maximizes his expected utility, i.e., he solves

$$\max_{a \in A} \sum_{\theta \in \Theta} \mu_m(\theta) u_R(a, \theta).$$

The maximum is well-defined since A is assumed to be a finite set. For convenience, define an auxiliary correspondence $\hat{a} : \Delta\Theta \rightrightarrows A$ by

$$\hat{a}(\mu) := \arg\max_{a \in A} \sum_{\theta \in \Theta} \mu(\theta) u_R(a, \theta).$$

In words, $\hat{a}(\mu)$ is the set of optimal actions of the receiver when his belief is μ.

The equilibrium concept we are considering is the Sender-preferred subgame perfect equilibrium. That is, whenever the receiver has multiple optimal actions under some belief, he chooses the one that maximizes the Sender's payoff. If such actions are still not unique, he chooses one of them using a "reasonable" tie-breaking rule.[4] There is some discussion on the robustness of choosing the Sender-preferred equilibrium, but in this chapter, we focus on the Sender-preferred subgame perfect equilibrium.

Definition 3.2. Suppose (M, π) is an information structure. Say *a message $m \in M$ leads to an action $a \in A$* if

$$a \in \hat{a}(\mu_m), \quad \text{and} \sum_{\theta \in \Theta} \mu_m(\theta) u_S(a, \theta) > \sum_{\theta \in \Theta} \mu_m(\theta) u_S(a', \theta), \ \forall a' \in \hat{a}(\mu_m) \setminus \{a\}.$$

[4] As shown in later proofs, the analysis in this chapter goes through as long as the tie-breaking rule for multiple Sender-preferred receiver-optimal action is rationalizable.

That is, m leads to a if a is the unique Sender-preferred receiver-optimal action under belief μ_m. Let $a^*(m)$ denote the action message m leads to.

Without further restriction, it is possible that two or more distinct actions in $\hat{a}(\mu_m)$ leads to the same expected payoff of the Sender. We make the following assumption to eliminate that possibility:

Assumption 3.2. For any belief $\mu \in \Delta\Theta$ and any pair of distinct actions $a_1, a_2 \in \hat{a}(\mu)$,

$$\sum_{\theta\in\Theta} \mu(\theta) u_S(a_1, \theta) \neq \sum_{\theta} \mu(\theta) u_S(a_2, \theta).$$

Assumption 3.2 guarantees that any message from any information structure leads to a unique action, or equivalently, the Sender-preferred Receiver-optimal action is unique for all belief μ. This is not so strong an assumption. For example, any persuasion problem with state independent Sender payoff $u_S : A \to \mathbb{R}$ and $u_S(a) \neq u_S(a')$ for any $a \neq a'$ satisfy this assumption.

With the necessary notations in place, the Sender's problem can be written as

$$\begin{aligned} \max_{(M,\pi)} \; & \sum_{m\in M} \left[\left(\sum_{\theta\in\Theta} \mu(\theta)\pi(m \mid \theta) \right) \left(\sum_{\theta\in\Theta} \mu_m(\theta) u_S(a^*(m), \theta) \right) \right] \\ \text{s.t.} \;\; & \mu_m \text{ and } a^*(m) \text{ are defined as above;} \\ & \pi(m \mid \theta) \in [\alpha, 1-\alpha], \; \forall m \in M, \; \forall \theta \in \Theta \end{aligned}$$

As shown in KG, the unconstrained problem (without the α-constraint) can be simplified as a belief-splitting problem. That is, instead of choosing information structures (M, π), it suffices to directly choose from all Bayes-plausible distribution over posteriors. This is sometimes referred to as "the belief approach." In our model, the belief approach is not directly applicable to general problems, due to the fact that the α-constraint adds new constraints on feasible distribution over posterior beliefs, i.e., not all Bayes-plausible distribution over posteriors are feasible under the α-constraint. Moreover, since the α-constraint is imposed on the signals, not the posteriors, it is not easy to have a clean translation of the α-constraint to an equivalent constraint on the posteriors. We will provide such a translation for the special case where both the state space and the action space are binary. More details are provided in Section 3.4.

3.3.1 Main Result

Since the belief-approach is not directly applicable, we turn to the "information-design" approach as called in Bergemann and Morris (2016). As it turns out, the "revelation principle" style result holds. That is, even in this constrained persuasion problem, the Sender gains nothing by using "complicated" message space instead of a message space as action recommendations to the Receiver.

Proposition 3.1 (Size of M or, "pseudo revelation principle."). *Fix a constrained persuasion problem* $(\Theta, \mu, A, u_S, u_R, \alpha)$. *Suppose it satisfies Assumptions 3.1 and 3.2. If an information structure* (π', M')

(i) is feasible under the α-constraint;

(ii) achieves persuasion payoff v; and

(iii) has $|M'| > |A|$.

Then there exists another information structure (π, M) that is feasible, achieves the same persuasion payoff v, and has $|M| = |A|$.

Proof. See Appendix C.1. □

The intuition behind Proposition 3.1 is clear. Whenever the Sender has $|M| > |A|$, two different messages necessarily lead to the same action of the Receiver. The

Sender can simply relabel this two messages by the action of the Receiver and add up their conditional probability as the new conditional probability. This operation is feasible under the α-constraint and preserves the Sender's persuasion payoff.

3.3.2 $|M| < |A|$ Could Be Uniquely Optimal

An important observation in the unconstrained persuasion problem is that having $M = A$ is always optimal. Even if some of the actions, say a, will never be recommended, the Sender can just use a signal (A, π) where $\pi(a \mid \theta) = 0$ for all states θ. Such a signal is not feasible under the α-constraint. In this section, we present a constrained persuasion problem $(\Theta, \mu, A, u_S, u_R, \alpha)$ in which having $|M| < |A|$, i.e., a message space as a proper subset of the action space, is uniquely optimal. Moreover, this example is not trivial in the sense that having $|M| = |A|$ is uniquely optimal in its corresponding unconstrained problem.

Example 3.1. Consider the following constrained persuasion problem with $\alpha = \frac{1}{6}$,

$$\Theta = \{-1, 0, 1\},\ \mu = (\mu_-, \mu_o, \mu_+) = \left(\frac{1}{4}, \frac{7}{12}, \frac{1}{6}\right), A = \{-1, 0, 1\}$$
$$u_S(a, \theta) = |a|,\ u_R(a, \theta) = -(a - \theta)^2.$$

To fix ideas, consider a hedge fund (Sender) trying to persuade a potential investor (Receiver) to invest in a stock. The investor wants to match the state, e.g., short-selling ($a = -1$) if the price is going down ($\theta = -1$), having no position ($a = 0$) if the price is not changing ($\theta = 0$), and buying in ($a = 1$) if the price is going up ($\theta = 1$). The hedge fund cares only about the commission that arises whenever the investor is acting ($a \neq 0$), regardless of the state. The common prior belief is as specified above.

First let's consider the unconstrained persuasion problem, Figure 3.2(a) below illustrates the function $v(\mu)$ that represents the Sender's highest payoff when the

Receiver's posterior is μ. Note that in Figure 3.2(b), the black dot on the bottom plane denotes the common prior $\mu = (\frac{1}{4}, \frac{7}{12}, \frac{1}{6})$ and it falls in the upper left corner triangle.

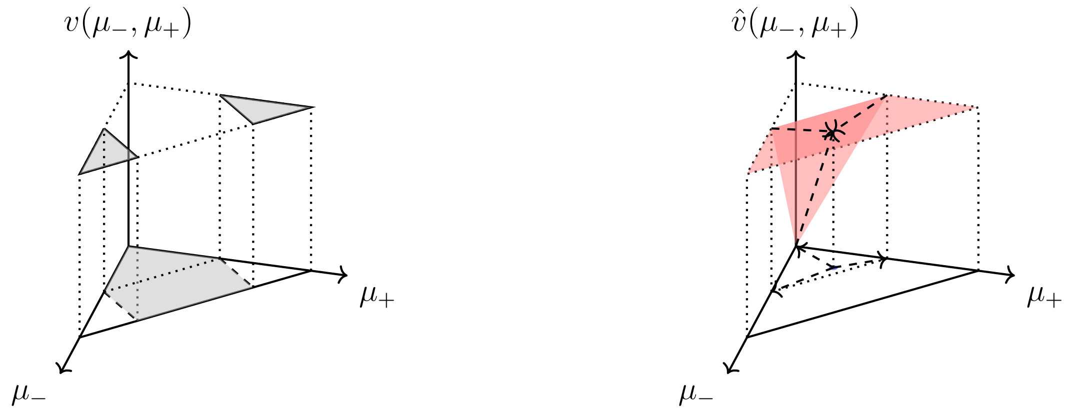

(a) Sender's payoff v as a function of μ when there is no persuasion

(b) Sender's optimal persuasion payoff $\hat{v}$, concavification of v.

Figure 3.2: Analysis for the unconstrained case, $\alpha = 0$.

When $\alpha = 0$, since the prior μ falls into the triangle region in the corner, the optimal information structure leads to the three posteriors $(0, \frac{1}{2}, \frac{1}{2})$, $(0, 1, 0)$ and $(\frac{1}{2}, \frac{1}{2}, 0)$. Having $|M| = 2$, i.e., only two possible posteriors, is suboptimal by applying the concavification result from KG and eyeballing Figure 3.2(b).

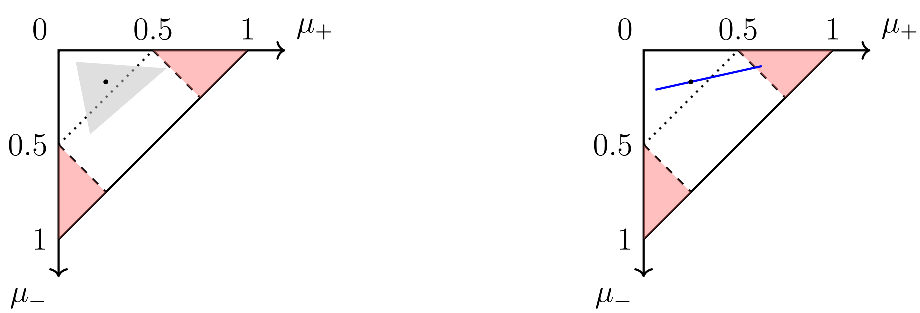

(a) $|M| = 3$, most informative signal no change in action

(b) $|M| = 2$, one signal change in action

Figure 3.3: Analysis for $\alpha = \frac{1}{6}$. $|M| = |A|$ is suboptimal.

When $\alpha = \frac{1}{6}$, however, having $|M| = |A| = 3$ is suboptimal. The reason is that with three messages, even the most informative information structure does not change the investor's mind enough to change his action. That is, having $|M| = |A| = 3$ leads to no improvement at all. But with $|M| = 2$, the most informative information structure does change the investor's mind enough to change his action, which leads to a strictly positive ex ante expected payoff for the hedge fund. Figure 3.3 gives a direct illustration of the intuition above.

3.3.3 An Algorithm to Solve the General Model

A natural implication from Proposition 3.1 is the following algorithm that can be used to solve any finite constrained persuasion problem.

Setup: collect all subsets of A that has two or more elements, i.e., a set $\mathcal{A}$ of sets, $\mathcal{A} := \{A_i\}_{i=1}^{2^{|A|}-|A|-1}$, where each A_i is a subset of A with $|A| \geq 2$ and $A_i \neq A_j$ if $i \neq j$.

Loop over i, i.e., loop over all these subsets of A. If A_i is picked out, use A_i as the set of actions that the Sender want to recommend, and solve the following problem:

$$\max_{\{\pi(a|\theta)\}_{a \in A_i,\ \theta \in \Theta}} \sum_{a \in A_i} \sum_{\theta \in \Theta} \mu(\theta)\pi(a_i \mid \theta) u_S(a_i, \theta)$$

$$\text{s.t.} \quad a_i \in \arg\max_{a' \in A_i} \sum_{\theta \in \Theta} \mu(\theta)\pi(a_i \mid \theta) u_R(a', \theta) \qquad \text{(obedience)}$$

$$\pi(a \mid \theta) \in [\alpha, 1-\alpha],\ \forall a \in A_i,\ \forall \theta \in \Theta \qquad \text{(feasibility)}$$

Note that the objective function and all constraints are linear in the choice variables. Hence this is a standard linear programming problem, and is easy to solve. Record the corresponding persuasion payoff (the objective function) and repeat until all sets in $\mathcal{A}$ are exhausted. Computationally speaking, although each subproblem is P-hard to solve, but the whole problem is NP-hard to solve, due to the fact that the number of subproblems is growing exponentially fast.

3.4 Belief-approach for Binary Action Space

Proposition 3.1 can be particularly useful when the action space is binary, i.e., $|A| = 2$, since we can safely conclude that the optimal signal must have binary message space, $|M| = 2$, no matter how large the state space Θ is. In this section, we will switch back to the belief-approach where we use distributions of posteriors as the choice variables when formulating the Sender's optimization problem. To do that, however, we need to characterize the feasible distributions over posteriors.

Definition 3.3 (Feasibility)**.** A collection of beliefs $(\mu_1, \ldots, \mu_n)$ (with $\mu_i \in \Delta\Theta$ for each i) is **feasible for the prior** μ if there exists a signal (M, π) such that:
(i) (M, π) is feasible under the α-constraint and
(ii) $(\mu_1, \ldots, \mu_n)$ is exactly the support of the distribution over posteriors induced by (M, π).

Without the α-constraint, a collection of beliefs $(\mu_1, \ldots, \mu_n)$ will be feasible as long as there exists some positive weights $(\lambda_1, \ldots, \lambda_n)$ that can average them back to the prior, i.e., $\sum_{i=1}^{n} \lambda_i \mu_i = \mu$ with $\sum_{i=1}^{n} \lambda_i = 1$. But with the α-constraint, that is no longer true.

3.4.1 Characterization of Feasible Pairs of Posteriors

For starters, consider a binary state space, $\Theta = \{\theta_1, \theta_2\}$. Let μ be the probability that θ_1 is the true state, and we can use $\mu \in [0, 1]$ to keep track of the belief over Θ.

With $|\Theta| = 2$ and $|M| = 2$, every signal $\pi : \Theta \to \Delta M$ can be summarized as a two-by-two matrix, where $p, q \in [\alpha, 1 - \alpha]$.

	θ_1	θ_2
m_1	p	q
m_2	$1 - p$	$1 - q$

Hence each π leads to a pair of posteriors (μ_1, μ_2) that can be reached with positive probability. Let μ_i denote the posterior belief when signal m_i is realized. Then $\mu_1 = \frac{p\mu}{p\mu + q(1-\mu)}$. Direct calculation gives that $\frac{\partial \mu_1}{\partial p} > 0$ and $\frac{\partial \mu_1}{\partial q} < 0$. Hence the posterior belief μ_1 is strictly increasing in p and strictly decreasing in q. With the α constraint, the upper bound of μ_1 for a given α and prior μ is given by letting $p = 1-\alpha$ and $q = \alpha$, i.e.,

$$\mathrm{UB}_\alpha(\mu) := \frac{(1-\alpha)\mu}{(1-\alpha)\mu + \alpha(1-\mu)} = \frac{(1-\alpha)\mu}{\alpha + (1-2\alpha)\mu}.$$

With similar arguments, the lower bound of μ_1 is

$$\mathrm{LB}_\alpha(\mu) := \frac{\alpha\mu}{\alpha\mu + (1-\alpha)(1-\mu)} = \frac{\alpha\mu}{1-\alpha + (2\alpha-1)\mu}.$$

Figure 3.4 is an illustration of these bounds as functions of the prior (with $\alpha = 0.1$).

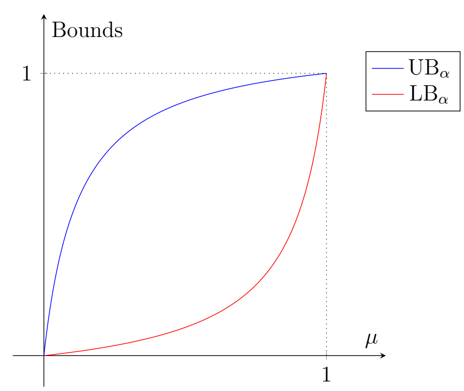

Figure 3.4: Bounds as functions of prior, $\alpha = 0.1$.

So for any given prior, one can first figure out the bounds. However, not every distribution over $[LB_\alpha(\mu), UB_\alpha(\mu)]$ that can average to μ is feasible. For example, the only way one can reach $LB_\alpha(\mu)$ is by choosing a message m_1 that has $\mathsf{P}(m_1 \mid \theta_1) = \alpha$ and $\mathsf{P}(m_1 \mid \theta_2) = 1-\alpha$, but this fixes the whole signal structure (due to the α-constraint), there can be exactly one other message m_2 with $\mathsf{P}(m_2 \mid \theta_1) = 1-\alpha$ and $\mathsf{P}(m_2 \mid \theta_2) = \alpha$. That is, if m_2 is sent, then the corresponding posterior is exactly $UB_\alpha(\mu)$. Therefore, if one specifies that $LB_\alpha(\mu)$ is one of the two posteriors that will be reached with positive probability, then the only pair of beliefs (μ_1, μ_2) with $\mu_1 = LB_\alpha(\mu)$ that can be the support of the distribution of posteriors through some signal structure is $(LB_\alpha(\mu), UB_\alpha(\mu))$.

Then it is natural to ask the following question: with $|M| = 2$, what are the feasible pairs of posterior beliefs (μ_1, μ_2) that can be generated as the support of the distribution of posteriors through some signal structure? Proposition 3.2 below provides an answer.

Proposition 3.2 (Characterization of feasible pairs of posteriors)**.** *Suppose* $\Theta = \{\theta_1, \theta_2\}$*, and the* α*-constraint is imposed. Fix any interior prior* μ*, then for any posterior* $\mu_1 \in [LB_\alpha(\mu), \mu)$*, the pair of beliefs* (μ_1, μ_2) *is feasible if and only if*

$$\mu_2 \in \left[\mu \cdot \frac{\mu_1 - \alpha\mu_1}{\mu_1 - \alpha\mu}, \frac{\mu - \mu_1 + \mu_1 \cdot \alpha(1-\mu)}{\mu - \mu_1 + \alpha(1-\mu)}\right],$$

for any $\mu_1 \in (\mu, UB_\alpha(\mu)]$*, the pair of beliefs* (μ_1, μ_2) *is feasible if and only if*

$$\mu_2 \in \left[\mu \cdot \frac{\alpha\mu_1}{\mu_1 - (1-\alpha)\mu}, \frac{\mu(1-\mu_1) - \mu_1 \cdot \alpha(1-\mu)}{1 - \mu_1 - \alpha(1-\mu)}\right].$$

Proof. See Appendix C.2. □

Figure 3.5 gives an illustration of the characterization given in Proposition 3.2. With $\alpha = 0.1$, and the prior belief $\mu = 0.4$. Every point (μ_1, μ_2) inside the shaded butterfly region is feasible, and every point other than (μ, μ) corresponds to a unique signal structure, while (μ, μ) can be reached through any completely uninformative signal (there are infinitely many of those). Every point outside the shaded butterfly region is not feasible.

When $\alpha = 0$ (the unconstrained standard case), the feasible region is the union of the rectangle on the northwest corner ($\mu_1 \leq \mu, \mu_2 \geq \mu$) and the rectangle on the southeast corner. An important feature of this unconstrained case is that the

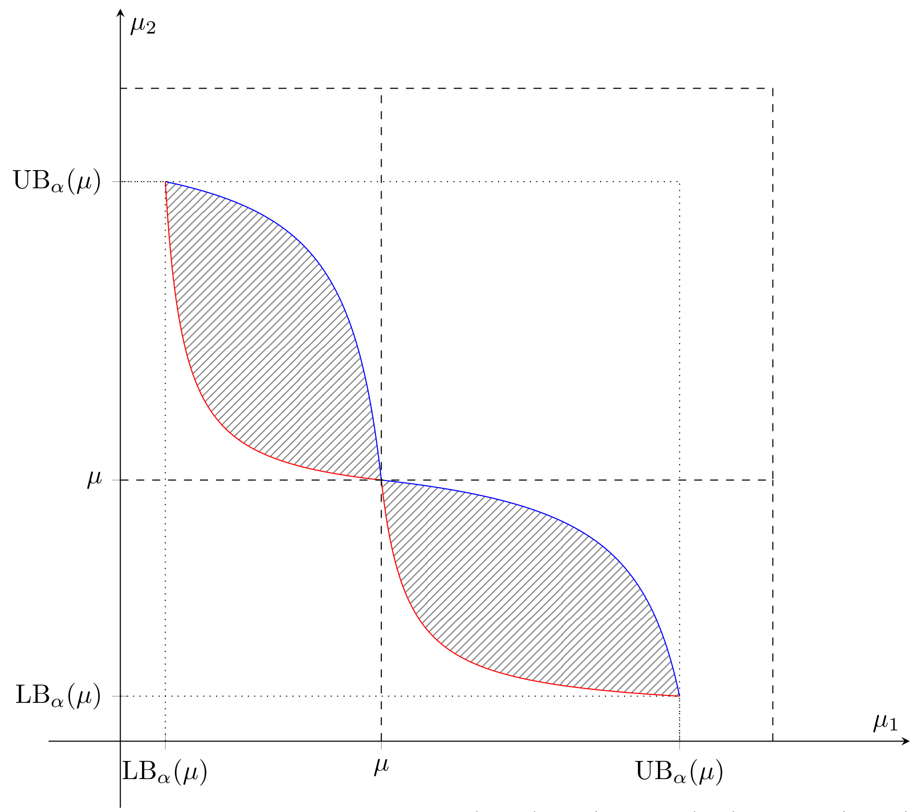

Figure 3.5: Feasible pairs of posteriors, $(\alpha, \mu) = (0.1, 0.4)$, $(\text{LB, UB}) = (\frac{2}{29}, \frac{6}{7})$.

"bounds" do not depend on any of the posterior. As long as one of the posterior $\mu_1 < \mu$ (or $\mu_1 > \mu$), any $\mu_2 \in (\mu, 1]$ (or $\mu_2 \in [0, \mu)$) will be feasible. That independence disappears when the α-constraint is introduced. As we can see from Figure 3.5, the "bounds" on feasible counterpart of one posterior changes with the other posterior.

Some more observation on Figure 3.5. The symmetry (about the 45 degree line) is intuitive, i.e., if a pair (μ_1, μ_2) is feasible, then so is (μ_2, μ_1). One can achieve this by simply switching the names of the signal realizations. It is also noteworthy that each half of the butterfly is convex, i.e., a linear combination of two feasible pairs is also feasible if their first coordinates are on the same side of μ.

To the best of our knowledge, it is an open question to extend the result of Proposition 3.2 to general cases where $|\Theta| > 2$, while keeping the assumption that $|A| = 2$.

3.5 Discussion

In this section, we will provide a comparison of the α-constraint considered in this chapter and other existing constraints on Sender's signal space.

3.5.1 α-constraint and Blackwell Informativeness Constraint

In Ichihashi (2019), the constraint on the Sender's signal space is modeled as following: there exists some pre-determined signal (M_*, π_*) that serves as an upper bound on how informative (in the sense of Blackwell (1953)) Sender's signal can be. The Sender can only use signals (M, π) that are less informative than (M_*, π_*).

We claim that comparing to the α-constraint, this is a weaker constraint (in the sense that it excludes fewer signals) when $|\Theta| = 2$, i.e., when the state space is binary, and these two constraints are not directly comparable when $|\Theta| \geq 3$.

Binary state space, $|\Theta| = 2$. To be clearer about our statements, define an auxiliary item Δ_α^n by

$$\Delta_\alpha^n := \left\{ (x_1, \ldots, x_m) \in \mathbb{R}^n \;\middle|\; \sum_{i=1}^{n} x_i = 1,\ x_i \in [\alpha, 1-\alpha],\ \forall i. \right\}$$

First we consider a binary state space, $|\Theta| = 2$. Write $\Theta = \{\theta_1, \theta_2\}$.

Intuitively, the signal (M, π) defined as below will be the "most informative" signal satisfying the α-constraint:

$(M, \pi):\ S = \{m_1, m_2\},\ \pi =$ (⋆)

	θ_1	θ_2
m_1	$1-\alpha$	α
m_2	α	$1-\alpha$

That intuition is actually correct. Formally, we have the following proposition:

Proposition 3.3. *Let (M, π) be defined as above. Suppose M' is any other message space with $|M'| = n \geq 2$. Let $\pi' : \Theta \to \Delta(M')$ be a signal such that $\pi'(\cdot \mid \theta) \in \Delta_\alpha^n$*

for all $\theta \in \{\theta_1, \theta_2\}$. *Then* (M, π) *is more informative than* (M', π') *in the Blackwell sense, i.e., there exists a garbling* $\gamma : S \to \Delta(S')$ *such that* $\pi' = \gamma \circ \pi$.

Proof. Suppose the garbling exists, then it can be described as the following matrix:

	m'_1	m'_2	$\cdots$	m'_n
m_1	x_1	x_2	$\cdots$	x_n
m_2	y_1	y_1	$\cdots$	y_n

Let $\pi'(\theta_1) = (p_1, \ldots, p_n)$ and $\pi'(\theta_2) = (q_1, \ldots, q_n)$. Then the problem transfers to solving the following system of equations (taking $\{p_i, q_i\}_{i=1}^n$ as given):

$$\begin{cases} p_i = (1-\alpha)x_i + \alpha y_i, \ \forall i = 1, \ldots, n \\ q_i = \alpha x_i + (1-\alpha) y_i, \ \forall i = 1, \ldots, n \\ \sum_i x_i = 1, \ \sum_i y_i = 1 \\ x_i \in [0,1], \ y_i \in [0,1], \ \forall i. \end{cases}$$

Using the first $2n$ equations, we get

$$\begin{cases} x_i = \dfrac{(1-\alpha)p_i - \alpha q_i}{1 - 2\alpha} \\ y_i = \dfrac{-\alpha p_i + (1-\alpha) q_i}{1 - 2\alpha} \end{cases}$$

By assumption, $p_i, q_i \in [\alpha, 1-\alpha]$ for all i. Hence $x_i, y_i \in [0,1]$. Moreover, with $\sum_i p_i = \sum_i q_i = 1$, it is indeed true that $\sum_i x_i = \sum_i y_i = 1$.

That is, we found a legitimate garbling γ satisfying $\pi' = \gamma \circ \pi$. □

One can have a more compact version of Proposition 3.3. Use $(M, \pi) \trianglerighteq (M', \pi')$ to denote that (M, π) is more informative than (M', π') in the Blackwell sense. Define

$$\Sigma_{(M,\pi)} := \Big\{ (M', \pi') \mid |M| < \infty, \ (M, \pi) \trianglerighteq (M', \pi') \Big\}.$$

That is, $\Sigma_{(M,\pi)}$ is the collection of all signals with finite message spaces (henceforth "finite signals") that are less informative than (M, π). In addition, define

$$\Sigma_\alpha := \Big\{ (M, \pi) \mid |M| < \infty, \ \pi(\cdot \mid \theta) \in \Delta_\alpha^{|S|}, \ \forall \theta \in \Theta \Big\}.$$

That is, Σ_α is the collection of all finite signals that satisfy the α-constraint.

Naturally, for fixed (M, π) and α, one can view these as correspondences that eats some set of states of the world Θ and spits out a set of available signals, so Proposition 3.3 can be re-written as

$$\Sigma_\alpha\big(\{\theta_1, \theta_2\}\big) \subseteq \Sigma_{(M,\pi)}\big(\{\theta_1, \theta_2\}\big), \text{ where } (M, \pi) \text{ is as specified in } (\star).$$

And the inclusion is strict since there are completely uninformative signals like

$(M', \pi') : \ M' = \{m_1, m_2\}, \ \pi' =$

	θ_1	θ_2
m_1	$1 - \frac{\alpha}{2}$	$1 - \frac{\alpha}{2}$
m_2	$\frac{\alpha}{2}$	$\frac{\alpha}{2}$

that are not in $\Sigma_\alpha(\cdot)$ but in $\Sigma_{(M,\pi)}(\cdot)$.

In this sense, the α-constraint is a stronger assumption than the Blackwell informativeness constraint to impose when Θ is binary.

General State Space with $|\Theta| \geq 3$. As the size of Θ grows from 2 to 3, things get much more complicated.

Let $\Theta := \{\theta_1, \theta_2, \theta_3\}$ and $\alpha < 1/4$. Consider the following potential "most informative signal" (M, π):

$(M, \pi) : \ M = \{m_1, m_2, m_3\}, \ \pi =$

	θ_1	θ_2	θ_3
m_1	$1 - 2\alpha$	α	α
m_2	α	$1 - 2\alpha$	α
m_3	α	α	$1 - 2\alpha$

This signal structure is nice in the following sense: it may be viewed as if the message is sent to the Receiver through a noisy channel where any m is sent out correctly as m with high probability $1-2\alpha$, while it gets send to the wrong messages $S - \{s\}$ with low probability α. However, the analogy of Proposition 3.3 is not true. That is, it is no longer true that every finite signal satisfying the α-constraint is less informative than (M, π) in the Blackwell sense.

To ease exposition, let $\alpha = 0.1$. Then (M, π) becomes

$(M, \pi): \; M = \{m_1, m_2, m_3\}, \; \pi =$

	θ_1	θ_2	θ_3
m_1	0.8	0.1	0.1
m_2	0.1	0.8	0.1
m_3	0.1	0.1	0.8

Consider the following signal structure (M', π'):

$(M', \pi'): \; M' = \{m'_1, m'_2\}, \; \pi' =$

	θ_1	θ_2	θ_3
m_1	0.9	0.1	0.1
m_2	0.1	0.9	0.9

Then straightforward calculation tells us that (M', π') as defined above satisfies the α-constraint but it is not comparable with (M, π) in terms of Blackwell informativeness. As seen in Example 3.1, this kinds of signals maybe invoked in the optimal persuasion mechanism. And there still exists completely uninformative signals that do not satisfy the α-constraint. Therefore, in this sense, the Blackwell informativeness constraint is not comparable with the α-constraint.

3.5.2 Connection with Other Approaches

The noisy channel approach. The noisy channel approach is closely related to the Blackwell informativeness approach in the following sense: If the communication channel is fixed, i.e., the input message space M and the output message space O are both fixed, and the matrix specifying the transition probabilities $p(o \mid m)$ (probability that an input message m results in an output message o) is fixed. Then this whole channel can be viewed as the upper bound signal, because any effective signal used by the Sender can be viewed as the signal chosen by the Sender composed with this channel.

The α-constraint then can be thought as a channel with an endogenous transition matrix, since the error rates are changing as the size of the message space is changing.

But the arguments to compare this two approaches would be similar to the arguments used in section 3.5.1.

The entropy reduction approach. The entropy reduction approach models the constraint by assuming that the reduction in expected entropy achieved by the Sender's signal cannot be larger than some threshold. And this constraint is easier to write using posterior beliefs. As argued above, the α-constraint is imposed directly on the signals, not the posteriors it induced. So it is hard to directly compare these two approaches. But as we see in Le Treust and Tomala (2019), the optimal persuasion subject to this entropy reduction constraint does not exhibit discountinuity in persuasion payoff as shown in the prosecutor judge example in Section 3.2. Moreover, the entropy reduction constraint will be slack if the optimal signal in the unconstrained case needs not be very informative. These observations seem to imply that the entropy reduction constraint is also different to the α-constraint.

3.6 Conclusion

This chapter studies constrained persuasion problems in which the Sender's signals need to satisfy an α-constraint. We find that the revelation principle style result in persuasion games continues to hold with the α-constraint, with the caveat that having a message space strictly smaller than the Receiver's actions space can be uniquely optimal. An algorithm is given to solve general constrained persuasion problems that fit into this model. We attempt to use the belief-approach as well, and provide a characterization of feasible posteriors in the special case where both the state space and the action space are binary.

Chapter 4

Conclusion

This dissertation contributes to the study of the role of information in individual decision-making and strategic interactions of economic agents. Chapter one contains the first model that directly links information avoidance with past choices and provides the natural explanation that economic agents might avoid information to reduce anticipated regret. Chapter two establishes several informativeness orders over ambiguous information structures as generalizations of Blackwell's celebrated notion of informativeness over precise information structures. Chapter three studies a constrained version of the Bayesian persuasion problem and compares the differences of the α-constraint with several existing constraints in the literature.

These chapters open doors for future research. The axiomatic nature of the model in chapter one provides the possibility for the model to be tested in lab experiments. It would also be fruitful to develop a more robust experimental environment for testing these more complicated behavioral implications. Chapter two leaves several open questions to be answered, among which the most interesting one would be how to establish an informativeness order over ambiguous information structures when the decision maker faces prior ambiguity. It would also be interesting to develop some axiomatic foundations for the monotone preference we have mainly focused on.

In this era of information, more and more observations and evidence on individuals' attitudes toward information emerge from everyday life. The author hopes to continue the study of individual preference for information to provide intuitive explanations for the observed behavioral patterns beyond the scope of this dissertation.

Appendix A

Appendix to Chapter 1

A.1 The Framework with Lotteries

In this section of the Appendix, we present some preliminary results in the framework with lotteries. These results will be used to prove results in the main text.

A.1.1 Lotteries

Let Z be a finite set of outcomes. Let $\Delta(Z)$ denote the set of lotteries over Z, with typical elements p, q. Endow $\Delta(Z)$ with the standard Euclidean metric. A menu is a nonempty compact subset of $\Delta(Z)$. Let $\widehat{\mathcal{M}}$ denote the set of all menus, with typical elements A, B. Endow $\widehat{\mathcal{M}}$ with the Hausdorff metric. A direction is a nonempty compact subset of $\widehat{\mathcal{M}}$. Let $\widehat{\mathcal{D}}$ denote the set of all directions, with typical elements $\mathbb{A}, \mathbb{B}$. Endow $\widehat{\mathcal{D}}$ with the Hausdorff metric.

Let $\mathcal{V}$ denote the set of normalized expected utilities over $\Delta(Z)$, that is,

$$\mathcal{V} := \left\{ v \in \mathbb{R}^Z \mid \sum_{z \in Z} v_z = 0 \right\}.$$

Let $\mathcal{U}$ denote the set of doubly normalized expected-utility functions on $\Delta(Z)$, that is,

$$\mathcal{U} := \left\{ u \in \mathbb{R}^Z : \sum_{z \in Z} u_z = 0, \sum_{z \in Z} u_z^2 = 1 \right\}.$$

A.1.2 Redundancy of a Collection of Utilities

An important notion that will be repeatedly used in our proof is about the redundancy of a collection of linear functions. Formally,

Definition A.1. Let $\{U_1, \dots, U_m\}$ be a collection of continuous linear functions from $\widehat{\mathcal{M}}$ to $\mathbb{R}$. We say this collection is *redundant* if there exists U_i that is constant or if there exists $i \neq j$ such that $U_j = \alpha U_i + \beta$ for some $\alpha > 0$ and $\beta \in \mathbb{R}$. We say a collection is *non-redundant* if it is not redundant.

Similarly, if $\{u_1, \dots, u_n\}$ is a collection of expected utility functions over $\Delta(Z)$, then we say this collection is *redundant* if there exists u_i that is constant or if there exists $i \neq j$ such that $u_j = \alpha u_i + \beta$ for some $\alpha > 0$ and $\beta \in \mathbb{R}$.

By convention, an empty collection is not redundant.

Using the notion of redundancy, we can prove the following lemma.

Lemma A.1. *Let $\{u_i\}_{i \in I}$ be a collection of normalized expected utility functions over $\Delta(Z)$, and let $U : \widehat{\mathcal{M}} \to \mathbb{R}$ be defined by $U(A) := \sum_{i \in I} \max_{p \in A} \; u_i(p)$. Then, $U(A) = 0$ for all $A \in \widehat{\mathcal{M}}$ if and only if $u_i(p) = 0$ for all $i \in I$ and all $p \in \Delta(Z)$.*

Proof of Lemma A.1. The "if" part is straightforward.

To prove the "only if" part, first let $J \subseteq I$ be a maximal collection of non-redundant subset of I, that is, for any $i \in I \setminus J$, u_i is constant or there exists some $j \in J$ such that $u_i = \alpha u_j$ for some $\alpha > 0$. It suffices to show that $U(A) \equiv 0$ implies $J = \emptyset$.

Suppose by contradiction that $|J| = 1$, then let $J = \{u_j\}$, and by the definition of J, there exists $L > 0$ such that $U(A) = L \max_{p \in A} u_j(p)$. And $L u_j(p) = U(\{p\}) = 0$ for all $p \in \Delta(Z)$ implies $u_j(p) = 0$, contradicting J being non-redundant. So $|J| \neq 1$.

Suppose by contradiction that $|J| \geq 2$, then $\{u_j\}_{j \in J}$ is a non-redundant collection of expected utilities, and by a standard result (e.g., Lemma A.1 of Kopylov (2009, JET)), there exists a menu of lotteries $A = \{p_j\}_{j\in J}$ such that p_j is the unique maximizer of u_j in A. Then, fix any $k \in J$,

$$U(A) = \sum_{j\in J} L_j \max_{p\in A} \; u_j(p) = \sum_{j\in J} L_j u_j(p_j) > \sum_{j\in J} L_j u_j(p_k) = U(\{p_k\}) = 0,$$

contradicting $U(A) = 0$. So $|J| < 2$. □

A.1.3 The Partial Regret Representation

Our primitive is a binary relation $\succsim$ over $\widehat{\mathcal{D}}$ (the set of all menus of menus of lotteries).

Definition A.2. A binary relation $\succsim$ over $\widehat{\mathcal{D}}$ has a *partial regret (PR) representation* if there exists a finitely-supported probability measure μ over $\mathcal{U}$ and a scalar $K \geq 0$ such that $\succsim$ is represented by the function $\widehat{V} : \widehat{\mathcal{D}} \to \mathbb{R}$ defined by

$$\widehat{V}(\mathbb{A}) = \max_{A\in\mathbb{A}} \sum_{u\in \mathrm{supp}(\mu)} \mu(u) \left[\max_{p\in A} u(p) - R(A, \mathbb{A}, u)\right] \tag{A.1}$$

where

$$R(A, \mathbb{A}, u) := K \left[\max_{B\in\mathbb{A}} \max_{p\in B} u(p) - \max_{p\in A} u(p)\right]. \tag{A.2}$$

The interpretation for the PR representation is the same as discussed in the main text. The agent has subjective uncertainty about her future tastes. This uncertainty resolves after her committing to a menu from a direction but before her choosing a lottery to consume from the menu. Given a taste, she can choose the best lottery from her menu of choice but inevitably suffers from regret if her choice of menu is suboptimal.

A useful equivalent expression of the utility representation is

$$\widehat{V}(\mathbb{A}) = \max_{A\in\mathbb{A}} \left[(1+K) \sum_{u\in \mathrm{supp}(\mu)} \mu(u) \max_{p\in A} u(p)\right] - K \sum_{u\in \mathrm{supp}(\mu)} \mu(u) \max_{B\in\mathbb{A}} \max_{p\in B} u(p). \tag{A.3}$$

We first present the axioms that characterize the PR representation. These should look familiar to the axioms presented in the main text. The proof for the representation theorem and some other intermediate results are in the next section.

Axiom A.1. (Weak Order): $\succsim$ *is complete and transitive.*

Axiom A.2. (Continuity): *For any* $\mathbb{A}$*, the sets* $\{\mathbb{B} : \mathbb{B} \succsim \mathbb{A}\}$, $\{\mathbb{B} : \mathbb{A} \succsim \mathbb{B}\}$ *are closed.*

Axiom A.3. (Independence): *If* $\mathbb{A} \succ \mathbb{B}$*, then for any* $\mathbb{C}$ *and* $\alpha \in (0, 1]$,

$$\alpha\mathbb{A} + (1 - \alpha)\mathbb{C} \succ \alpha\mathbb{B} + (1 - \alpha)\mathbb{C}.$$

For the next axiom, we need to define the notion of a critical subset in the framework with lotteries. The interpretation is very similar to that in the framework with acts.

Definition A.3. Let $\mathbb{A}$ be a direction. Say that $\mathbb{B}$ is *critical for* $\mathbb{A}$ if $\mathbb{B} \subseteq \mathbb{A}$ and $\mathbb{B}' \sim A$ for all B' satisfying $\mathbb{B} \subseteq \mathbb{B}' \subseteq \mathbb{A}$. Say that B *is critical for* A *in* $\mathbb{A}$ if $B \subseteq A \in \mathbb{A}$ and $(\mathbb{A} \setminus \{A\}) \cup \{B'\} \sim \mathbb{A}$ for all B' satisfying $B \subseteq B' \subseteq A$.

Axiom A.4. (Finiteness): *There exists a natural number* N *such that*

- *For every* $\mathbb{A} \in \mathcal{D}$*, there exists* $\mathbb{B}$ *with* $|\mathbb{B}| < N$ *such that* $\mathbb{B}$ *is critical for* $\mathbb{A}$ *;*
- *For every* $\mathbb{A} \in \mathcal{D}$ *and every* $A \in \mathbb{A}$*, there exists* B *with* $|B| < N$ *such that* B *is critical for* A *in* $\mathbb{A}$.

Axiom A.5. (Ex-Ante Regret): *If* $\{A\} \succsim \{B\}$ *and* $A \in \mathbb{A}$*, then* $\mathbb{A} \succsim \mathbb{A} \cup \{B\}$.

Axiom A.6. (Interim Preference for Flexibility): *For any direction* $\mathbb{A}$ *and any menus* A, B, $\mathbb{A} \cup \{A \cup B\} \succsim \mathbb{A} \cup \{A, B\}$.

Axiom A.7. (Inclusion): *If* $B \subseteq A$ *and* $A \in \mathbb{A}$*, then* $\mathbb{A} \cup \{B\} \succsim \mathbb{A}$.

Axiom A.8. (Nontriviality): *There exists* $\mathbb{A}$ *and* $\mathbb{B}$ *such that* $\mathbb{B} \subseteq \mathbb{A}$ *with* $\mathbb{A} \succ \mathbb{B}$*.*

Axioms A.1-A.5 corresponds to Axioms 1.1-1.5, respectively. The inclusion axiom corresponds to a weakened version of the domination axiom (Axiom 1.8), and the nontriviality axiom is stated slightly different than Axiom 1.7.

Theorem A.2. *A binary relation* $\succsim$ *over* $\widehat{\mathcal{D}}$ *has a partial regret representation if and only if it satisfies Axioms A.1-A.8.*

A.1.4 Proof of Theorem A.2

We start by establishing a nested DLR representation from Axioms A.1-A.4. DLR refers to Dekel, Lipman, and Rustichini (2009) who establishes a finite version of Dekel, Lipman, and Rustichini (2001) for representations with subjective state spaces.

Definition A.4. A binary relation $\succsim$ over $\widehat{\mathcal{D}}$ has a *nested DLR representation* if $\succsim$ can be represented by

$$V_{DLR}(\mathbb{A}) = \sum_{i \in I^+} \max_{A \in \mathbb{A}} U_i(A) - \sum_{i \in I^-} \max_{A \in \mathbb{A}} U_i(A) \tag{A.4}$$

where I^+ is the index set for the positive states and I^- is the index set for the negative states and $\{U_i\}_{i \in I^+ \cup I^-}$ is a non-redundant collection of continuous linear functions from $\widehat{\mathcal{M}}$ to $\mathbb{R}$. (It is without loss to assume that $I^+ \cap I^- = \emptyset$. Let $I := I^+ \cup I^-$.) Moreover, for each $i \in I$, $U_i : \widehat{\mathcal{M}} \to \mathbb{R}$ has a DLR representation with

$$U_i(A) = \sum_{k \in P_i} \max_{p \in A} w_{ik}(p) - \sum_{j \in N_i} \max_{p \in A} v_{ij}(p) \tag{A.5}$$

where P_i is the index set for the positive substates for U_i and N_i is the index set of negative substates for U_i (it is without loss to assume that $P_i \cap N_i = \emptyset$) and $\{w_k\}_{k \in P_i} \cup \{v_j\}_{j \in N_i}$ is a non-redundant collection of normalized expected utilities over $\Delta(Z)$.

Lemma A.3. *A binary relation* $\succsim$ *over* $\widehat{\mathcal{D}}$ *has a nested DLR representation if and only if it satisfies Axioms A.1-A.4.*

Proof. See the proof of Theorem 5 of Stovall (2018). □

Given the nested DLR representation, we can start to fine-tune the states and substates through the other axioms to get to the PR representation.

Axiom A.9. (Strong Ex-Ante Regret): *If for any* $B \in \mathbb{B}$*, there exists* $A \in \mathbb{A}$ *such that* $\{A\} \succsim \{B\}$*, then* $\mathbb{A} \succsim \mathbb{A} \cup \mathbb{B}$*.*

As suggested by its name, Axiom A.9 is a strengthening of Axiom A.5. However, together with Axioms A.1 and A.2, Axioms A.5 and A.9 are equivalent. This is summarized in the following lemma.

Lemma A.4. *Suppose* $\succsim$ *satisfies Axioms A.1 (Weak Order) and A.2 (Continuity), then it satisfies Axiom A.5 if and only if it satisfies Axiom A.9.*

Proof. It is straightforward to see that if $\succsim$ satisfies Axiom A.9, then it satisfies Axiom A.5. We want to show that the other direction goes through, that is, Axiom A.5 implies Axiom A.9. Suppose $\succsim$ satisfies Axiom A.5, and suppose two directions of lotteries $\mathbb{A}$ and $\mathbb{B}$ are such that for all $B \in \mathbb{B}$, there exists $A \in \mathbb{A}$ such that $\{A\} \succsim \{B\}$.

First note that the set of all menus of lotteries, $\widehat{\mathcal{M}}$, is compact when equipped with the Hausdorff metric. Therefore, $\widehat{\mathcal{M}}$ is separable. We can thus choose a countable dense subset of $\mathbb{B}$, say $\mathbb{B}^* = \{B_i \mid i = 1, 2, \dots\}$. Define a sequence of directions $(\mathbb{B}_n)_{n=1,2,\dots}$ by $\mathbb{B}_n = \{B_1, B_2, \dots, B_n\}$. Then $\mathbb{B}_n \subseteq \mathbb{B}_{n+1}$ for all n and $\mathbb{B}^* = \bigcup_{n=1}^{\infty} \mathbb{B}_n$. By a standard result, $(\mathbb{B}_n)_{n=1,2,\dots}$ converges to the closure of $\mathbb{B}^*$, that is, $\mathbb{B}_n \to \overline{\mathbb{B}^*} = \mathbb{B}$ as n goes to infinity. Note that $B_1 \in \mathbb{B}^* \subseteq \mathbb{B}$, so there exists $A_1 \in \mathbb{A}$ such that $\{A_1\} \succsim \{B_1\}$ by assumption, and Axiom A.5 implies that $\mathbb{A} \succsim \mathbb{A} \cup \{B_1\}$. Let

$\mathbb{A}_1 = \mathbb{A} \cup \{B_1\} = \mathbb{A} \cup \mathbb{B}_1$. Similarly, $B_2 \in \mathbb{B}^* \subseteq \mathbb{B}$, so there exists $A_2 \in \mathbb{A}$ such that $\{A_2\} \succsim \{B_2\}$, and Axiom A.5 implies that $\mathbb{A}_1 \succsim \mathbb{A}_1 \cup \{B_2\}$. Let $\mathbb{A}_2 = \mathbb{A}_1 \cup \{B_2\} = \mathbb{A} \cup \mathbb{B}_2$. By transitivity, $\mathbb{A} \succsim \mathbb{A}_1$ and $\mathbb{A}_1 \succsim \mathbb{A}_2$ imply that $\mathbb{A} \succsim \mathbb{A}_2 = \mathbb{A} \cup \mathbb{B}_2$. We can repeat this argument for each n and by an induction argument, $\mathbb{A} \succsim \mathbb{A} \cup \mathbb{B}_n$ for all n. Therefore, $\mathbb{A} \cup \mathbb{B}_n$ is in the lower contour set of $\mathbb{A}$ for all n, that is, $\mathbb{A} \cup \mathbb{B}_n \in \{\mathbb{C} \in \widehat{\mathcal{D}} \mid \mathbb{A} \succsim \mathbb{C}\}$. Since $\succsim$ satisfies Axiom A.2 (Continuity), this lower contour set is closed, $\mathbb{A} \succsim \lim_{n\to\infty} \mathbb{A} \cup \mathbb{B}_n = \mathbb{A} \cup \mathbb{B}$. □

With this equivalence established, we show the effect of imposing Axiom A.9 on a nested DLR representation.

Lemma A.5. *A binary relation $\succsim$ over $\widehat{\mathcal{D}}$ satisfies Axioms A.1-A.4 and Axiom A.9 if and only if the following two conditions hold:*

1. *$\succsim$ has a nested DLR representation with at most one positive state, that is, $|I^+| \leq 1$ in equation (A.4); and*

2. *Let $\widehat{v}(A) := V_{DLR}(\{A\}) = \sum_{i\in I^+} U_i(A) - \sum_{i\in I^-} U_i(A)$. There exists a non-negative scalar $\alpha \geq 0$ such that $\sum_{i\in I^+} U_i(A) = \alpha\widehat{v}(A)$ for all $A \in \widehat{\mathcal{M}}$.*

Proof of Lemma A.5.

If. Condition 1 implies that $\succsim$ has a nested DLR representation, so by Lemma A.3, $\succsim$ satisfies Axioms A.1-A.4. We just need to check if Axiom A.9 is implied by conditions 1 and 2. If $|I^+| = 0$, that is, $I^+ = \emptyset$, then there is no positive state. Thus, $\mathbb{A} \succsim \mathbb{A} \cup \mathbb{B}$ for any two directions $\mathbb{A}, \mathbb{B} \in \widehat{\mathcal{D}}$. Axiom A.9 is satisfied.

If $|I^+| = 1$, then there is exactly one positive state, and

$$V_{DLR}(\mathbb{A}) = \max_{A\in\mathbb{A}} U_0(A) - \sum_{i\in I^-} \max_{A\in\mathbb{A}} U_i(A),$$

with $\{U_0\} \cup \{U_i\}_{i\in I^-}$ being a non-redundant collection of continuous linear functions from $\widehat{\mathcal{M}}$ to $\mathbb{R}$. Recall that $\widehat{v}(A) := V_{DLR}(\{A\}) = \sum_{i\in I^+} U_i(A) - \sum_{i\in I^-} U_i(A)$, and $\{A\} \succsim \{B\}$ if and only if $\widehat{v}(A) \geq \widehat{v}(B)$.

Suppose for any $B \in \mathbb{B}$, there exists $A \in \mathbb{A}$ such that $\widehat{v}(A) \geq \widehat{v}(B)$. Since U_0 is non-constant, condition 2 implies that $\alpha > 0$, and

$$U_0(A) = \alpha \underbrace{\left[U_0(A) - \sum_{i\in I^-} U_i(A)\right]}_{=\widehat{v}(A)=V_{DLR}(\{A\})} \text{ for all } A \in \widehat{\mathcal{M}}.$$

Therefore, condition 2 implies that for any $B \in \mathbb{B}$, there exists $A \in \mathbb{A}$ such that $U_0(A) \geq U_0(B)$, which further indicates that $\max_{A\in\mathbb{A}\cup\mathbb{B}} U_0(A) = \max_{A\in\mathbb{A}} U_0(A)$. Since $\mathbb{A}\cup\mathbb{B} \supseteq \mathbb{A}$, $\max_{A\in\mathbb{A}\cup\mathbb{B}} U_i(A) \geq \max_{A\in\mathbb{A}} U_i(A)$ for any $i \in I^-$. Thus, $V_{DLR}(\mathbb{A}) \geq V_{DLR}(\mathbb{A}\cup\mathbb{B})$, and $\mathbb{A} \succsim \mathbb{A}\cup\mathbb{B}$. Axiom A.9 is satisfied when $|I^+| = 1$.

This completes the proof of the "if" part.

Only if. By Lemma A.3, Axioms A.1-A.4 imply the existence of a nested DLR representation for $\succsim$ as in Definition A.4. We want to show that if Axiom A.9 is also satisfied, then $|I^+| \leq 1$ and condition 2 is also satisfied.

If $|I^+| = 0$, then $I^+ = \emptyset$ and there is no positive state, condition 1 is satisfied. And we can set $\alpha = 0$ so that condition 2 is satisfied.

If $|I^+| > 0$, then $I^+ \neq \emptyset$ and there is at least one positive state. It suffices to show that: If Axiom A.9 is satisfied, then for any positive state $i \in I^+$, there exists $\alpha > 0$ such that $U_i(A) = \alpha\widehat{v}(A)$. Suppose by contradiction that for some $i^* \in I^+$, U_{i^*} does not represent the same preference over $\widehat{\mathcal{M}}$ as $\widehat{v}$. Consider three cases:

- Case 1: $\widehat{v}$ represents a trivial preference over $\widehat{\mathcal{M}}$, that is, $\widehat{v}(A) = 0$ for all $A \in \widehat{\mathcal{M}}$. Then there is at least another state in $I^+ \cup I^-$ that is different from i^* (otherwise $\widehat{v}(A) = U_{i^*}(A)$, contradiction). So there are at least two elements in $\{U_k\}_{k\in I^+\cup I^-}$.

Thus we can apply a standard result (e.g., Lemma A.1 of Kopylov, 2009) to conclude that there exists a collection of menus $\mathbb{A} := \{A_k\}_{k \in I^+ \cup I^-}$ such that A_k is the unique maximizer of U_k in $\mathbb{A}$ for each $k \in I^+ \cup I^-$, and $|\mathbb{A}| \geq 2$.

Let $\mathbb{A}_1 := \mathbb{A} \setminus \{A_{i^*}\}$ and $\mathbb{B}_1 := \{A_{i^*}\}$. Then $\mathbb{A}_1 \neq \emptyset$ (since $|\mathbb{A}| \geq 2$) and for any $B \in \mathbb{B}_1$, there exists $A \in \mathbb{A}_1$ such that $\{A\} \succsim \{B\}$ (since $\widehat{v}(A) = 0 = \widehat{v}(B)$ for any $A, B \in \widehat{\mathcal{M}}$). Moreover, $\max_{A \in \mathbb{A}_1 \cup \mathbb{B}_1} U_{i^*}(A) > \max_{A \in \mathbb{A}_1} U_{i^*}(A)$ and $\max_{A \in \mathbb{A}_1 \cup \mathbb{B}_1} U_k(A) = \max_{A \in \mathbb{A}_1} U_k(A)$ for any $k \neq i^*$. Therefore, $V_{DLR}(\mathbb{A}_1 \cup \mathbb{B}_1) > V_{DLR}(\mathbb{A}_1)$, indicating $\mathbb{A}_1 \cup \mathbb{B}_1 \succ \mathbb{A}_1$, violating Axiom A.9.

- Case 2: $\{\widehat{v}\} \cup \{U_k\}_{k \in I^+ \cup I^-}$ is non-redundant.

  There are at least two elements ($\widehat{v}$ and U_{i^*}) in this collection, thus we can apply a standard result (e.g., Lemma A.1 of Kopylov, 2009) to conclude that there exists a collection of menus $\mathbb{A}' := \{B\} \cup \{A_k\}_{k \in I^+ \cup I^-}$ such that B is unique maximizer of $\widehat{v}$ in $\mathbb{A}$ and A_k is the unique maximizer of U_k in $\mathbb{A}$ for each $k \in I^+ \cup I^-$.

  Construct two directions by $\mathbb{A}'_1 := \mathbb{A}' \setminus \{A_{i^*}\}$ and $\mathbb{B}'_1 := \{A_{i^*}\}$. Then $\mathbb{A}'_1 \neq \emptyset$ (since $B \in \mathbb{A}'_1$) and $\mathbb{A}'_1 \geq \mathbb{B}'_1$ (since $B \in \mathbb{A}'_1$ and $\widehat{v}(B) > \widehat{v}(A_{i^*})$). Moreover, $\max_{A \in \mathbb{A}'_1 \cup \mathbb{B}'_1} U_{i^*}(A) > \max_{A \in \mathbb{A}'_1} U_{i^*}(A)$ and $\max_{A \in \mathbb{A}'_1 \cup \mathbb{B}'_1} U_k(A) = \max_{A \in \mathbb{A}'_1} U_k(A)$ for any $k \neq i^*$. Therefore, $V_{DLR}(\mathbb{A}'_1 \cup \mathbb{B}'_1) > V_{DLR}(\mathbb{A}'_1)$, indicating $\mathbb{A}'_1 \cup \mathbb{B}'_1 \succ \mathbb{A}'_1$, violating Axiom A.9.

- Case 3: There exists some $j \in I^+ \cup I^-$ such that $j \neq i^*$ but $\widehat{v}$ represents the same preference over $\widehat{\mathcal{M}}$ as U_j. Then $\{\widehat{v}\} \cup \{U_k\}_{k \in I^+ \cup I^- \setminus \{j\}}$ is non-redundant. Then the similar arguments as in Case 2 above will lead to a violation of Axiom A.9.

This completes the proof of the "only if" part. □

We move on to further fine-tune the states and substates of the nested DLR representation. A weapon for that is the following lemma.

Lemma A.6. *Let $\{U_i\}_{i \in I}$ be a finite collection of continuous linear functions from $\widehat{\mathcal{M}}$ to $\mathbb{R}$ such that the collection is non-redundant, and each $U_i : \widehat{\mathcal{M}} \to \mathbb{R}$ has a minimal finite DLR representation, that is,*

$$U_i(A) = \sum_{k \in P_i} \max_{p \in A} w_{ik}(p) - \sum_{j \in N_i} \max_{p \in A} v_{ij}(p)$$

where for each $i \in I$, $\{w_{ik}\}_{k \in P_i} \cup \{v_{ij}\}_{j \in N_i}$ is a non-redundant collection of normalized expected utilities on $\Delta(Z)$. Then there exists a collection of menus $\{A_i\}_{i \in I}$ (all in the interior of $\Delta(Z)$) such that:

1. *$U_i(A_i) > U_i(A_j)$ for any $i \in I$ and any $j \neq i$.*

2. *For any $i \in I$, any $k \in P_i$ and any $j \in N_i$,*

$$\left| \arg\max_{p \in A_i} w_{ik}(p) \right| = 1, \quad \left| \arg\max_{q \in A_i} v_{ij}(q) \right| = 1,$$

   and for each fixed i, all these unique maximizers are distinct from each other.

Proof. See the proof of Lemma 3 of Stovall (2018). □

With Lemma A.6 in hand, we can further fine-tune the substates in the nested DLR representation.

The next lemma states that imposing the Interim Preference for Flexibility axiom is equivalent to requiring there to be no negative substates in any positive state and at most one positive substate in any negative state. Formally,

Lemma A.7. *Suppose $\succsim$ has a nested DLR representation as in Definition A.4, then $\succsim$ satisfies Axiom A.6 (Interim Preference for Flexibility) if and only if $|N_i| = 0$ for all $i \in I^+$ and $|P_i| \leq 1$ for all $i \in I^-$.*

Proof. Suppose $\succsim$ has a nested DLR representation as in Definition A.4, moreover, $|N_i| = 0$ for any $i \in I^+$ and $|P_i| \leq 1$ for any $i \in I^-$. We want to show that $\mathbb{A} \cup \{A \cup B\} \succsim \mathbb{A} \cup \{A, B\}$ for any direction $\mathbb{A}$ and menus A, B.

For any positive state $i \in I^+$, $U_i(A \cup B) \geq \max\{U_i(A), U_i(B)\}$ since $|N_i| = 0$. Therefore, the positive terms in $V_{DLR}(\mathbb{A} \cup \{A \cup B\})$ is weakly larger than the positive terms in $V_{DLR}(\mathbb{A} \cup \{A, B\})$.

For any negative state $i \in I^-$: If $|P_i| = 0$, then $U_i(C) = -\sum_{j \in N_i} \max_{p \in C} v_{ij}(p)$, and $-U_i(A \cup B) \geq \max\{-U_i(A), -U_i(B)\}$. If $|P_i| = 1$, then

$$U_i(C) = \max_{p \in C} w_{i0}(p) - \sum_{j \in N_i} \max_{p \in C} v_{ij}(p),$$

and since $\max_{p \in A \cup B} w_{i0}(p) = \max\{\max_{p \in A} w_{i0}(p), \max_{p \in B} w_{i0}(p)\}$, we can again conclude that $-U_i(A \cup B) \geq \max\{-U_i(A), -U_i(B)\}$. Therefore, the negative terms in $V_{DLR}(\mathbb{A} \cup \{A \cup B\})$ is weakly larger than the negative terms in $V_{DLR}(\mathbb{A} \cup \{A, B\})$.

This completes the proof for the "if" part.

Only if. Suppose $\succsim$ has a nested DLR representation as in Definition A.4, moreover, $\succsim$ satisfies Axiom A.6 (Interim Preference for Flexibility). We want to show that $|N_i| = 0$ for all $i \in I^+$ and $|P_i| \leq 1$ for all $i \in I^-$.

By assumption, $\succsim$ can be represented by

$$V_{DLR}(\mathbb{A}) = \sum_{i \in I^+} \max_{A \in \mathbb{A}} U_i(A) - \sum_{i \in I^-} \max_{A \in \mathbb{A}} U_i(A)$$

where $\{U_i\}_{i \in I^+ \cup I^-}$ is a non-redundant collection of continuous linear functions from $\widehat{\mathcal{M}}$ to $\mathbb{R}$. Let $I := I^+ \cup I^-$. For each $i \in I$,

$$U_i(A) = \sum_{k \in P_i} \max_{p \in A} w_{ik}(p) - \sum_{j \in N_i} \max_{p \in A} v_{ij}(p),$$

where $\{w_{ik}\}_{k \in P_i} \cup \{v_{ij}\}_{j \in N_i}$ is a non-redundant collection of normalized expected-utilities on $\Delta(Z)$. Thus, all assumptions of Lemma A.6 are satisfied. Therefore, there exists a collection of menus $\{A_i\}_{i\in I}$ in the interior of $\Delta(Z)$ such that

1. $U_i(A_i) > U_i(A_j)$ for any $i, j \in I$ with $j \neq i$.

2. For any $i \in I$, any $k \in P_i$ and any $j \in N_i$,

$$\left|\arg\max_{p\in A_i} w_{ik}(p)\right| = 1, \ \left|\arg\max_{q\in A_i} v_{ij}(q)\right| = 1,$$

   and for each fixed i, all these unique maximizers are distinct from each other.

Let $\mathbb{A} = \{A_i\}_{i\in I}$.

We will first show that $|P_i| \leq 1$ for all $i \in I^-$, that is, there is at most one positive sub-state in any negative state. We do this by proving its contrapositive.

Suppose by contradiction that $|P_{i^*}| \geq 2$ for some $i^* \in I^-$, that is, U_{i^*} has two or more positive sub-states. Then we want to construct A, B such that

$$V_{DLR}(\mathbb{A} \cup \{A \cup B\}) < V_{DLR}(\mathbb{A} \cup \{A, B\}).$$

We want A, B and $A \cup B$ to be all "closed" to A_{i^*} so that for any $i \in I$ with $i \neq i^*$,

$$U_i(A_i) > U_i(A_{i^*}) \approx \max\Big\{U_i(A \cup B), U_i(A), U_i(B)\Big\}$$

This will guarantee that the difference between $V_{DLR}(\mathbb{A} \cup \{A \cup B\})$ and $V_{DLR}(\mathbb{A} \cup \{A, B\})$ is generated solely by the difference in U_{i^*}. Since i^* is a negative state, to get the desired strict inequality, we want

$$\max_{C\in\mathbb{A}\cup\{A\cup B\}} U_{i^*}(C) > \max_{C\in\mathbb{A}\cup\{A,B\}} U_{i^*}(C).$$

For ease of exposition, we will write U_* for U_{i^*} and A_* for A_{i^*}. Similarly, we write P_* for P_{i^*} (the index set of positive sub-states in i^*), N_* for N_{i^*}, w_{*k} for w_{i^*k} for any $k \in P_*$ and v_{*j} for v_{i^*j} for any $j \in N_*$. For $a, b \in \mathbb{R}$, write $a \vee b$ to denote $\max\{a, b\}$.

By construction, A_* is the unique maximizer of U_* in $\mathbb{A}$, thus, it suffices to have

$$U_*(A_*) \vee U_*(A \cup B) > U_*(A_*) \vee U_*(A) \vee U_*(B)$$

which will happen if we have $U_*(A \cup B) > U_*(A_*), U_*(A), U_*(B)$. With this goal in mind, we start to construct A and B.

Define $p_k := \arg\max_{p \in A_*} w_{*k}(p)$. Since $|P_*| \geq 2$, we can pick k and k' with $k \neq k'$ from P_*. For any $\varepsilon > 0$, define

$$A^\varepsilon := A_* \cup \{p_k + \varepsilon w_{*k}\},\ B^\varepsilon := A_* \cup \{p_{k'} + \varepsilon w_{*k'}\}.$$

Since w_{*k} is normalized and p_k is in the interior of $\Delta(Z)$ by construction, by a standard result, we can find some ε small enough so that $p_k + \varepsilon w_{*k}$ is also in the interior of $\Delta(Z)$.

Note that for any $\varepsilon > 0$,

$$w_{*k}(p_k + \varepsilon w_{*k}) = w_{*k}(p_k) + \varepsilon \|w_{*k}\|^2 > w_{*k}(p_k).$$

Therefore,

$$\max_{p \in A^\varepsilon} w_{*k}(p) = w_{*k}(p_k) \vee w_{*k}(p_k + \varepsilon w_{*k}) = w_{*k}(p_k) + \varepsilon \|w_{*k}\|^2$$

$$\max_{p \in B^\varepsilon} w_{*k}(p) = w_{*k}(p_k) \vee \Big(w_{*k}(p_{k'}) + \varepsilon w_{*k}(w_{*k'})\Big)$$

By construction, $w_{*k}(p_{k'}) < w_{*k}(p_k)$, so there exists $\varepsilon_1 > 0$ such that

$$w_{*k}(p_{k'}) + \varepsilon_1 w_{*k}(w_{*k'}) < w_{*k}(p_k),$$

which further indicates that $\max_{p \in B^{\varepsilon_1}} w_{*k}(p) = w_{*k}(p_k)$. With similar arguments, we could find some $\varepsilon_2 > 0$ such that

$$\max_{p \in B^{\varepsilon_2}} w_{*k'}(p) > w_{*k'}(p_{k'}) = \max_{p \in A^{\varepsilon_2}} w_{*k'}(p).$$

Let $\varepsilon_3 := \min\{\varepsilon_1, \varepsilon_2\}$, then

$$\max_{p \in A^{\varepsilon_3}} w_{*k}(p) > w_{*k}(p_k) = \max_{p \in B^{\varepsilon_3}} w_{*k}(p)$$

$$\max_{p \in A^{\varepsilon_3}} w_{*k'}(p) = w_{*k'}(p_{k'}) < \max_{p \in B^{\varepsilon_3}} w_{*k'}(p).$$

Note that

$$U_*(A^{\varepsilon_3} \cup B^{\varepsilon_3}) = \sum_{k \in P_*} \max_{p \in A^{\varepsilon_3} \cup B^{\varepsilon_3}} w_{*k}(p) - \sum_{j \in N_*} \max_{p \in A^{\varepsilon_3} \cup B^{\varepsilon_3}} v_{*j}(p)$$

$$U_*(A^{\varepsilon_3}) = \sum_{k \in P_*} \max_{p \in A^{\varepsilon_3}} w_{*k}(p) - \sum_{j \in N_*} \max_{p \in A^{\varepsilon_3}} v_{*j}(p)$$

$$U_*(B^{\varepsilon_3}) = \sum_{k \in P_*} \max_{p \in B^{\varepsilon_3}} w_{*k}(p) - \sum_{j \in N_*} \max_{p \in B^{\varepsilon_3}} v_{*j}(p)$$

For each $j \in N_*$,

$$\max_{p \in A^{\varepsilon_3} \cup B^{\varepsilon_3}} v_{*j}(p) = \max_{p \in A_1} v_{*j}(p) \vee v_{*j}(p_k + \varepsilon_3 w_{*k}) \vee v_{*j}(p_{k'} + \varepsilon_3 w_{*k'}).$$

Since A_* is finite and v_{*j} has a unique maximizer in A_* that is not p_k or $p_{k'}$ (by Lemma*), we have a small number $\varepsilon_j > 0$ with $\varepsilon_j < \varepsilon_3$ such that

$$\max_{p \in A^{\varepsilon_j} \cup B^{\varepsilon_j}} v_{*j}(p) = \max_{p \in A^{\varepsilon_j}} v_{*j}(p) = \max_{p \in B^{\varepsilon_j}} v_{*j}(p) = \max_{p \in A_*} v_{*j}(p).$$

Let $\varepsilon := \min_{j \in N_*} \varepsilon_j$, then A^ε and B^ε will satisfy our desired condition:

$$U_*(A^\varepsilon \cup B^\varepsilon) > U_*(A_*), U_*(A^\varepsilon), U_*(B^\varepsilon)$$

because: (i) they all have the same negative term, and (ii) $A^\varepsilon \cup B^\varepsilon \supset A^\varepsilon, B^\varepsilon \supseteq A_*$ so $A^\varepsilon \cup B^\varepsilon$ has a weakly larger positive term, but it must be strictly larger as well because A^ε contains a unique maximizer $p_k + \varepsilon w_{*k}$ in sub-state k and B^ε contains a unique maximizer $p_{k'} + \varepsilon w_{*k'}$ in sub-state k'.

Finally, we can further scale down ε to guarantee that for any $i \in I$ with $i \neq i^*$,

$$U_i(A_i) > U_i(A^\varepsilon \cup B^\varepsilon), U_i(A^\varepsilon), U_i(B^\varepsilon)$$

because: (i) A_i is the unique maximizer of U_i in $\mathbb{A}$ and (ii) $A^\varepsilon \cup B^\varepsilon$, A^ε and B^ε can all be made arbitrarily close to A_*.

We will then show that for any $i \in I^+$, $|N_i| = 0$, that is, there is no negative sub-state in any positive state. We do this by proving its contrapositive.

Suppose by contradiction that $|N_{i^*}| \geq 1$ for some $i^* \in I^+$. Since $|N_*| \geq 1$, we can fix some $j \in N_*$. Let $q_j := \arg\max_{p \in A_*} v_{*j}(p)$. For any $\varepsilon > 0$, define

$$A^\varepsilon := A_* \quad \text{and} \quad B^\varepsilon = (A_* \setminus \{q_j\}) \cup \{q_j - \varepsilon v_{*j}\}.$$

We can find ε small enough such that $q_j - \varepsilon v_{*j}$ is in the interior of $\Delta(Z)$. For any such ε, $v_{*j}(q_j - \varepsilon v_{*j}) = v_{*j}(q_j) - \varepsilon \|v_{*j}\|^2 < v_{*j}(q_j) = \max_{p \in A_*} v_{*j}(p)$. Moreover, we can make ε small enough so that

$$v_{*j}(q_j - \varepsilon v_{*j}) = v_{*j}(q_j) - \varepsilon \|v_{*j}\|^2 > \max_{p \in A_* \setminus \{q_j\}} v_{*j}(p)$$

$$v_{*j'}(q_j - \varepsilon v_{*j}) = v_{*j'}(q_j) - \varepsilon v_{*j'}(v_{*j}) < \max_{p \in A_*} v_{*j'}(p), \ \forall j' \in N_* \text{ and } j' \neq j$$

$$w_{*k}(q_j - \varepsilon v_{*j}) = w_{*k}(q_j) - \varepsilon w_{*k}(v_{*j}) < \max_{p \in A_*} w_{*k}(p), \ \forall k \in P_*$$

Fix such a ε, we will have

$$U_*(B^\varepsilon) > U_*(A^\varepsilon) = U_*(A_*) = U_*(A^\varepsilon \cup B^\varepsilon).$$

We can further scale down ε so that $U_i(A_i) > U_i(A^\varepsilon), U_i(B^\varepsilon)$ for any $i \in I$ and $i \neq i^*$. Then $V_{DLR}(\mathbb{A} \cup \{A^\varepsilon, B^\varepsilon\}) > V_{DLR}(\mathbb{A} \cup \{A^\varepsilon \cup B^\varepsilon\})$, violating Axiom A.6.

This completes the proof of Lemma A.7. $\square$

Lemma A.8 states that imposing the Inclusion axiom is equivalent to requiring there to be no negative substates in any negative state. Formally,

Lemma A.8. *Suppose $\succsim$ has a nested DLR representation as in Definition A.4, then $\succsim$ satisfies Axiom A.7 (Inclusion) if and only if $|N_i| = 0$ for any $i \in I^-$.*

Proof. If. Suppose $\succsim$ has a nested DLR representation as in Definition A.4, moreover, $|N_i| = 0$ for any $i \in I^-$. We want to show that $\mathbb{A} \cup \{B\} \succsim \mathbb{A}$ for any direction $\mathbb{A}$ and any menu B such that $B \subseteq A$ for some $A \in \mathbb{A}$.

Fix A, B with $B \subseteq A$. For any negative state $i \in I^-$, $U_i(A) \geq U_i(B)$ since $|N_i| = 0$. Since $A \in \mathbb{A}$, we must have $-\max_{C \in \mathbb{A}} U_i(C) = -\max_{C \in \mathbb{A} \cup \{B\}} U_i(C)$.

For any positive state $i \in I^+$, $\max_{C\in\mathbb{A}\cup\{B\}} U_i(C) \geq \max_{C\in\mathbb{A}} U_i(C)$. Thus, $V_{DLR}(\mathbb{A} \cup \{B\}) \geq V_{DLR}(\mathbb{A})$.

Only if. Suppose $\succsim$ has a nested DLR representation as in Definition A.4, moreover, $\succsim$ satisfies Axiom A.7 (Inclusion) is satisfied, we want to show that for any $i \in I^-$, $|N_i| = 0$. That is, there is no negative sub-state in any negative state. We do this by proving its contrapositive.

Let $\mathbb{A} = \{A_i\}_{i\in I}$ be constructed the same way as in the proof of Lemma A.7. Suppose by contradiction that $|N_{i^*}| \geq 1$ for some $i^* \in I^-$. For ease of exposition, write U_* for U_{i^*}, A_* for A_{i^*}, P_* for P_{i^*}, N_* for N_{i^*}, w_{*k} for w_{i^*k} for any $k \in P_*$, and v_{*j} for v_{i^*j} for any $j \in N_*$.

Since $|N_*| \geq 1$, we can fix some $j \in N_*$. Let $q_j := \arg\max_{p\in A_*} v_{*j}(p)$. For any $\varepsilon > 0$, define

$$A^\varepsilon := A_* \cup \{q_j - \varepsilon v_{*j}\} \quad \text{and} \quad B^\varepsilon = (A_* \setminus \{q_j\}) \cup \{q_j - \varepsilon v_{*j}\}.$$

We can find ε small enough such that $q_j - \varepsilon v_{*j}$ is in the interior of $\Delta(Z)$. For any such ε, $v_{*j}(q_j - \varepsilon v_{*j}) = v_{*j}(q_j) - \varepsilon\|v_{*j}\|^2 < v_{*j}(q_j) = \max_{p\in A_*} v_{*j}(p)$. Moreover, we can make ε small enough so that

$$v_{*j}(q_j - \varepsilon v_{*j}) = v_{*j}(q_j) - \varepsilon\|v_{*j}\|^2 > \max_{p\in A_*\setminus\{q_j\}} v_{*j}(p)$$

$$v_{*j'}(q_j - \varepsilon v_{*j}) = v_{*j'}(q_j) - \varepsilon v_{*j'}(v_{*j}) < \max_{p\in A_*} v_{*j'}(p),\ \forall j' \in N_* \text{ and } j' \neq j$$

$$w_{*k}(q_j - \varepsilon v_{*j}) = w_{*k}(q_j) - \varepsilon w_{*k}(v_{*j}) < \max_{p\in A_*} w_{*k}(p),\ \forall k \in P_*$$

Fix such a ε, we will have

$$U_*(B^\varepsilon) > U_*(A^\varepsilon) = U_*(A_*).$$

We can further scale down ε so that $U_i(A_i) > U_i(A^\varepsilon), U_i(B^\varepsilon)$ for any $i \in I$ and $i \neq i^*$. Then $V_{DLR}(\mathbb{A} \cup \{A^\varepsilon, B^\varepsilon\}) < V_{DLR}(\mathbb{A} \cup \{A^\varepsilon\})$ (since i^* is a negative state), which is a direct violation of Inclusion. $\square$

Now we are ready to present the proof for Theorem A.2.

Proof of Theorem A.2.

If. Suppose $\succsim$ has a FPR representation with parameters (μ, K), that is, $\succsim$ can be represented by

$$\widehat{V}(\mathbb{A}) = \max_{A\in\mathbb{A}}\left[(1+K)\sum_{u\in\text{supp}(\mu)}\mu(u)\max_{p\in A}u(p)\right] - K\sum_{u\in\text{supp}(\mu)}\mu(u)\max_{A\in\mathbb{A}}\max_{p\in A}u(p).$$

We want to show that Axioms A.1-A.8 are satisfied. It is straightforward to verify the necessity of the axioms using Lemmas A.3, A.5, A.7, A.8.

Only if. Axioms A.1-A.4 delivers a nested DLR representation (Lemma A.3). Axiom A.5 guarantees that there is at most one positive state (part 1 of Lemma A.5). Axiom A.6 guarantees that there are no negative substates in any positive state and at most one positive substate in any negative state (Lemma A.7). Axiom A.7 guarantees that there are no negetive substates in any negative state (Lemma A.8). Thus, there can be at most one positive state (which contains only positive substates) and there might be multiple negative states (each of which contains at most one positive substate and no negative substates).

Therefore, Axioms A.1-A.7 imply that there exists two non-redundant collections of normalized expected utilities over $\Delta(Z)$, $\{w_k\}_{k\in P}$ and $\{v_j\}_{j\in N}$, such that $\succsim$ can be represented by the function $V_S : \widehat{\mathcal{D}} \to \mathbb{R}$ defined by

$$V_S(\mathbb{A}) = \max_{A\in\mathbb{A}}\sum_{k\in P}\max_{p\in A}w_k(p) - \sum_{j\in N}\max_{A\in\mathbb{A}}\max_{p\in A}v_j(p). \tag{A.6}$$

For convenience, define

$$U_0(A) := \sum_{k\in P}\max_{p\in A}w_k(p) \qquad U_j(A) := \max_{p\in A}v_j(p) \text{ for each } j\in N$$

$$\widehat{v}(A) := V_S(\{A\}) = \sum_{k\in P}\max_{p\in A}w_k(p) - \sum_{j\in N}\max_{p\in A}v_j(p).$$

If $\succsim$ can be represented by V_S as defined in equation (A.6) and $\succsim$ satisfies Axiom A.8, then $|P| \geq 1$ and U_0 is non-constant. Otherwise there will be no positive state in the representation and $\mathbb{A} \succsim \mathbb{B}$ for any $\mathbb{A} \subseteq \mathbb{B}$, violating Axiom A.8.

Note that since $\{v_j\}_{j \in N}$ is non-redundant, $\{U_j\}_{j \in N}$ must also be non-redundant. We continue our proof by proving the following lemma.

Lemma A.9. *If $\succsim$ can be represented by V_S defined in equation* (A.6) *with* $|P| \geq 1$ *and $\succsim$ satisfies Axiom A.5, then there exists $\alpha > 0$ such that $U_0 = \alpha v$.*

Proof of Lemma A.9. We discuss two possible cases.

Case 1: Suppose $\{U_0\} \cup \{U_j\}_{j \in N}$ is non-redundant. Then

$$V_S(\mathbb{A}) := \max_{A \in \mathbb{A}} U_0(A) - \sum_{j \in N} \max_{A \in \mathbb{A}} U_j(A)$$

is a nested DLR representation with exactly one positive state. By part 2 of Lemma A.5, this implies that there exists $\alpha > 0$ such that $U_0 = \alpha v$.

Case 2: Suppose $\{U_0\} \cup \{U_j\}_{j \in N}$ is redundant. By our previous results, $\{U_j\}_{j \in N}$ is not redundant and U_0 is non-constant. Therefore, there exists some $\beta > 0$ and exactly one $j \in N$ such that $U_0 = \beta U_j$. (N cannot be empty, otherwise U_0 is the only function in the collection and $\{U_0\}$ is not redundant.) It must be that $\beta > 1$, otherwise there will be no positive state (U_0 is absorbed by U_j) and Axiom A.8 will be violated.

$U_0 = \beta U_j$ means that for any menu A,

$$\sum_{k \in P} \max_{p \in A} w_k(p) = \beta \max_{p \in A} v_j(p).$$

Now this can only happen when $|P| = 1$, otherwise the LHS will exhibit strict preference for flexibility in some cases while the RHS will always exhibit strategic rationality. Let $P = \{k\}$, then $w_k = \beta v_j$, and $U_0 = \beta \max_{p\in A} v_j(p)$. Let $U_0' := (\beta - 1)U_j$, then $\{U_0'\} \cup \{U_{j'}\}_{j'\in N,\ j'\neq j}$ is non-redundant. Then we can apply Lemma A.5 again to conclude that there exists $\alpha > 0$ such that $U_0' = \alpha\widehat{v}$ where

$$\widehat{v}(A) := V_S(\{A\}) = U_0(A) - \sum_{j\in N} U_j(A) = U_0'(A) - \sum_{j'\in N,\ j'\neq j} U_{j'}(A).$$

Therefore,

$$U_0 = \frac{\beta}{\beta-1}U_0' = \frac{\beta}{\beta-1}(\alpha\widehat{v}).$$

With $\beta > 1$ and $\alpha > 0$, we can set $\alpha' := \frac{\alpha\beta}{\beta-1}$. Then $\alpha' > 0$ and $U_0 = \alpha'\widehat{v}$.

This completes the proof of Lemma A.9. □

We have just shown that if Axioms A.1-A.8 are satisfied, then there exists $\alpha > 0$ such that $U_0 = \alpha\widehat{v}$. That is,

$$\sum_{k\in P} \max_{p\in A} w_k(p) = \alpha\left[\sum_{k\in P} \max_{p\in A} w_k(p) - \sum_{j\in N} \max_{p\in A} v_j(p)\right]$$

which further indicates that

$$\sum_{j\in N} \max_{p\in A} v_j(p) = \frac{\alpha-1}{\alpha} \sum_{k\in P} \max_{p\in A} w_k(p). \tag{A.7}$$

Claim: For equation (A.7) to hold, it must be that $\alpha \geq 1$.

Proof of the Claim. Suppose by contradiction that $\alpha < 1$. Then $(\alpha-1)/\alpha < 0$, and the RHS of equation (A.7) will represent a preference that exhibits the opposite of preference for flexibility, while the LHS of equation (A.7) represents a preference that exhibits preference for flexibility, contradiction. □

Now if $\alpha = 1$, then for any menu A,

$$\sum_{j\in N} \max_{p\in A} v_j(p) = 0.$$

By Lemma A.1, this indicates that $N = \emptyset$. And

$$V_S(\mathbb{A}) = \max_{A \in \mathbb{A}} \sum_{k \in P} \max_{p \in A} w_k(p)$$

with $|P| \geq 1$. We can then get a FPR representation by setting $K = 0$ and (doubly) normalizing each w_k.

If $\alpha > 1$, then $|N| \geq 1$, and for any menu A,

$$U_0(A) = \sum_{k \in P} \max_{p \in A} w_k(p) = \frac{\alpha}{\alpha - 1} \sum_{j \in N} \max_{p \in A} v_j(p),$$

which further indicates that

$$V_S(\mathbb{A}) = \max_{A \in \mathbb{A}} \left[\frac{\alpha}{\alpha - 1} \sum_{j \in N} \max_{p \in A} v_j(p) \right] - \sum_{j \in N} \max_{A \in \mathbb{A}} \max_{p \in A} v_j(p).$$

Since $\alpha > 1$, we can get a FPR representation by setting $K = \alpha - 1 > 0$, (doubly) normalizing each v_j and scaling everything up by multiplicating K.

This completes the proof of Theorem A.2. □

A.1.5 Uniqueness of the PR representation

For the identification of the parameters μ and K, we build on the identification result of Dekel et al. (2001) and the uniqueness results in Sarver (2008).

Let $\widehat{V}$ be a PR representation for $\succsim$ with parameters (μ, K).

Define $\widehat{v} : \widehat{\mathcal{M}} \to \mathbb{R}$, for any menu $A \in \widehat{\mathcal{M}}$, by

$$\widehat{v}(A) := \widehat{V}(\{A\}) = \sum_{u \in \text{supp}(\mu)} \mu(u) \max_{p \in A} u(p). \tag{A.8}$$

That is, $\widehat{v}$ represents a preference over $\widehat{\mathcal{M}}$ generated by $\succsim$ restricting to singleton directions (directions containing only one menu). This is a special case of the representation captured in Dekel, Lipman, and Rustichini (2001) and Dekel, Lipman, Rustichini, and Sarver (2007) with preference for flexibility. Therefore, we can apply their identification result to conclude that the subjective belief over tastes, μ, is uniquely identified (since the expected utilities are doubly normalized). Formally,

Lemma A.10. *Suppose both (μ, K) and (μ', K') represent $\succsim$, then $\mu' = \mu$.*

Now that μ is identified, we move on to the identification of K. We follow a similar chain of steps used to establishing the uniqueness results in Sarver (2008).

Define $\widehat{r} : \widehat{\mathcal{D}} \to \mathbb{R}$, for any direction $\mathbb{A} \in \widehat{\mathcal{D}}$, by

$$\begin{aligned}\widehat{r}(\mathbb{A}) &:= \min_{A \in \mathbb{A}} \sum_{u \in \text{supp}(\mu)} \mu(u) R(A, \mathbb{A}, u) \\ &= \min_{A \in \mathbb{A}} \sum_{u \in \text{supp}(\mu)} \mu(u) K \left[\max_{B \in \mathbb{A}} \max_{q \in B} u(q) - \max_{p \in A} u(p) \right].\end{aligned} \tag{A.9}$$

The function $\widehat{r}(\mathbb{A})$ represents the minimal expected regret that the agent can experience when faced with direction $\mathbb{A}$. Note that given a direction $\mathbb{A}$, the menu A that maximizes v also minimizes expected regret. Therefore, for any direction $\mathbb{A} \in \widehat{\mathcal{D}}$, the agent will choose $A \in \mathbb{A}$ to maximize $v(A)$, and

$$\widehat{V}(\mathbb{A}) = \max_{A \in \mathbb{A}} \widehat{v}(A) - \widehat{r}(\mathbb{A}). \tag{A.10}$$

Theorem A.11. *Two PR representations (μ, K) and (μ', K') represent the same preference $\succsim$ if and only if $\widehat{v}' = \widehat{v}$ and $\widehat{r}' = \widehat{r}$.*

Proof of Theorem A.11. The "if" part is straightforward.

For the "only if" part, suppose (μ, K) and (μ', K') represent the same preference, then we can first apply the mixture space theorem to guarantee that there exists $\alpha > 0$ such that $\widehat{v}' = \alpha \widehat{v}$ and $\widehat{r}' = \alpha \widehat{r}$.

But by Lemma A.10, $\mu' = \mu$ implies that $\widehat{v}' = \widehat{v}$. Thus, $\alpha = 1$, and $\widehat{r}' = \widehat{r}$. $\square$

To proceed with the identification of K, we first rule out a less interesting case where Axiom A.5 is trivially satisfied.

Lemma A.12. *Suppose $\succsim$ has a PR representation with parameters (μ, K), then the following are equivalent:*

1. $\widehat{r}(\mathbb{A}) = 0$ *for all* $\mathbb{A} \in \widehat{\mathcal{D}}$*;*

2. $\succsim$ *satisfies monotonicity. That is, if* $\mathbb{B} \subseteq \mathbb{A}$*, then* $\mathbb{A} \succsim \mathbb{B}$*.;*

3. $K = 0$ *or* $\mu = \delta_u$ *for some* $u \in \mathcal{U}$ *(or both).*

Proof of Lemma A.12. We show that $1 \iff 2$ and $1 \iff 3$.

$1 \Longrightarrow 2$: Suppose $\widehat{r}(\mathbb{A}) = 0$ for all $\mathbb{A} \in \widehat{\mathcal{D}}$. If $\mathbb{A} \supseteq \mathbb{B}$, then $\widehat{V}(\mathbb{A}) = \max_{A \in \mathbb{A}} \widehat{v}(A) \geq \max_{B \in \mathbb{B}} \widehat{v}(B) = \widehat{V}(\mathbb{B})$. Thus, $\mathbb{A} \succsim \mathbb{B}$.

$2 \Longrightarrow 1$: Suppose by contradiction that $\widehat{r}(\mathbb{A}) > 0$ for some $\mathbb{A} \in \widehat{\mathcal{D}}$, we want to show that $\succsim$ will not satisfy monotonicity. Fix a direction $\mathbb{A}$ such that $\widehat{r}(\mathbb{A}) > 0$, let $A \in \arg\max_{B \in \mathbb{A}} \widehat{v}(B)$, then $\{A\} \subseteq \mathbb{A}$ but $\{A\} \succ \mathbb{A}$ since

$$\widehat{V}(\{A\}) = \widehat{v}(A) > \widehat{v}(A) - \widehat{r}(\mathbb{A}) = \widehat{V}(\mathbb{A}).$$

Taking the contrapositive completes the proof.

$3 \Longrightarrow 1$: Straightforward.

$1 \Longrightarrow 3$: Suppose by contradiction that $K > 0$ and $|\mathrm{supp}(\mu)| \geq 2$, then there exists a menu of lotteries $A_0 = \{p_u\}_{u \in \mathrm{supp}(\mu)}$ such that p_u is the unique maximizer of u in A. Let $\mathbb{A} := \{\{p_u\} : u \in \mathrm{supp}(\mu)\}$. Then

$$\begin{aligned}
\widehat{r}(\mathbb{A}) &= \min_{A \in \mathbb{A}} \sum_{u \in \mathrm{supp}(\mu)} \mu(u) K \left[\max_{B \in \mathbb{A}} \max_{q \in B} u(q) - \max_{p \in A} u(p) \right] \\
&= \min_{p \in A_0} \sum_{u \in \mathrm{supp}(\mu)} \mu(u) K \left[\max_{q \in A_0} u(q) - u(p) \right]
\end{aligned}$$

But for any $p_u \in A_0$, $\max_{q \in A_0} u'(q) - u'(p) > 0$ for any $u' \neq u$. So $\widehat{r}(\mathbb{A}) > 0$. Taking the contrapositive completes the proof. □

We say $\succsim$ has a nontrivial PR representation if there exist $\mathbb{A}$ and $\mathbb{B}$ such that $\mathbb{B} \subseteq \mathbb{A}$ but $\mathbb{B} \succ \mathbb{A}$.

Theorem A.13. *Suppose $\succsim$ has a nontrivial PR representation and $\succsim$ has two PR representations (μ, K) and (μ', K'), then $\mu' = \mu$ and $K' = K$.*

Proof. $\mu' = \mu$ by Lemma A.10, and by Theorem A.11, $\widehat{v}' = \widehat{v}$, $\widehat{r}' = \widehat{r}$. Since $\succsim$ has a nontrivial FPR representation, $K, K' > 0$. Since $\mu' = \mu$, it must be that $\widehat{r}(\mathbb{A})/K = \widehat{r}'(\mathbb{A})/K'$ for all $\mathbb{A} \in \widehat{\mathcal{D}}$. Thus, $\widehat{r}(\mathbb{A}) = \widehat{r}'(\mathbb{A})$ for all $\mathbb{A}$ implies that $K' = K$. □

A.2 Proof of Theorem 1.1

A.2.1 Necessity of Axioms 1.1-1.8

Suppose $\succsim$ has a SIT representation with parameters (π, u, K, σ), we want to show that $\succsim$ satisfies Axioms 1.1-1.8.

By assumption, $\succsim$ can be represented by

$$V(\mathbb{F}) = \max_{F \in \mathbb{F}} \left[(1+K) \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega) u(f(\omega)) \right]$$
$$- K \sum_{s \in S} \max_{G \in \mathbb{F}} \max_{g \in G} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega) u(g(\omega))$$

It is without loss to assume that $0 \notin S$. For convenience, let

$$U_0(F) := \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega) u(f(\omega)) \tag{A.11}$$

$$U_s(F) := \max_{g \in G} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega) u(g(\omega)) \text{ for each } s \in S \tag{A.12}$$

It can be easily verified that U_0 and $(U_s)_{s \in S}$ are continuous linear functions from $\mathcal{M}$ to $\mathbb{R}$, where $\mathcal{M}$ is the set of all menus of acts equppied with the Hausdorff metric.

Therefore, the binary relation $\succsim$ can be represented by

$$V(\mathbb{F}) = \max_{F\in\mathbb{F}}(1+K)U_0(F) - \sum_{s\in S}\max_{F\in\mathbb{F}} K\cdot U_s(F). \tag{A.13}$$

This is a finite DLR type representation over a general convex space, which is studied in Kopylov (2009).

Lemma A.14. *If $\succsim$ has a SIT representation, then $\succsim$ satisfies Axioms 1.1-1.3 (Weak Order, Continuity and Independence).*

Proof of Lemma A.14. By the arguments above, equation (A.13) implies that $\succsim$ has a finite DLR type representation over a general domain as characterized in Kopylov (2009). Therefore, $\succsim$ satisfies Axioms 1 and 2 by applying Theorem 2.1 of Kopylov (2009).

Axiom 1.3, our independence axiom, is slightly stronger than the independence axiom posited in Kopylov (2009). To be precise, we verify that $\succsim$ satisfies Axiom 1.3 directly. Let $\mathbb{F},\mathbb{G},\mathbb{H}$ be three directions and $\alpha\in(0,1)$.

$$\begin{aligned}
\mathbb{F}\succsim\mathbb{G} &\iff V(\mathbb{F})\geq V(\mathbb{G})\\
&\iff \alpha V(\mathbb{F})+(1-\alpha)V(\mathbb{H})\geq \alpha V(\mathbb{F})+(1-\alpha)V(\mathbb{H})\\
&\iff V\big(\alpha\mathbb{F}+(1-\alpha)\mathbb{H}\big)\geq V\big(\alpha\mathbb{G}+(1-\alpha)\mathbb{H}\big)\\
&\iff \alpha\mathbb{F}+(1-\alpha)\mathbb{H}\succsim\alpha\mathbb{G}+(1-\alpha)\mathbb{H}
\end{aligned}$$

where the third equivalence follows from the definition of the convex combination of two directions and the fact that $U_0,(U_s)_{s\in S}$ are all linear functions satisfying

$$U_s(\alpha F+(1-\alpha)G)=\alpha U_s(F)+(1-\alpha)U_s(G)$$

for any menus $F,G\in\mathcal{M}$ and any scalar $\alpha\in[0,1]$. □

Lemma A.15. *If $\succsim$ has a SIT representation, then $\succsim$ satisfies Axioms 1.4 (Finiteness).*

Proof of Lemma A.15. Recall that S is a finite set representing the possible signal realizations from the anticipated information structure σ. Let $N = |S| + 2$.

For any direction $\mathbb{F}$ such that $|\mathbb{F}| < N$, the direction itself is a critical subset whose cardinality is smaller than N. For any direction $\mathbb{F}$ such that $|\mathbb{F}| \geq N$ (including the countable and uncountable cases), let $F_0 \in \arg\max_{F \in \mathbb{F}} U_0(F)$, and $F_s \in \arg\max_{F \in \mathbb{F}} U_s(F)$ for each $s \in S$. Then $\mathbb{G} := \{F_0\} \cup \{F_s\}_{s \in S}$ is critical for $\mathbb{F}$ and $|\mathbb{G}| \leq |S| + 1 < N$. For the second part of Axiom 4, fix any direction $\mathbb{F}$ and menu $F \in \mathbb{F}$. If $|F| < N$, then F itself is critical for F in $\mathbb{F}$. If $F \notin \arg\max_{G \in \mathbb{F}} U_0(G)$ and $F \notin \arg\max_{G \in \mathbb{F}} U_s(G)$ for any $s \in S$, then this menu does not matter in the first place for the evaluation of $\mathbb{F}$, so the empty set $\emptyset$ is critical for F in $\mathbb{F}$. Lastly, if $|F| \geq N$ and F is a maximizer of some of U_0 and $(U_s)_{s \in S}$, we discuss two different cases:

- If F is a maximizer of U_0, then let f_s be a maximizer of $\sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega)u\big(f(\omega)\big)$ for each $s \in S$, then $G := \{f_s\}_{s \in S}$ is critical for F in $\mathbb{F}$, and $|G| \leq |S| < N$.

- If F is a maximizer of U_s for some $s \in S$, let g be a maximizer of $\sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega)u\big(f(\omega)\big)$, then $G := \{g\}$ is critical for F in $\mathbb{F}$, and $|G| = 1 < N$.

This completes the proof of Lemma A.15. $\square$

Lemma A.16. *If $\succsim$ has a SIT representation, then $\succsim$ satisfies Axioms 1.5 (Ex-Ante Regret).*

Proof of Lemma A.16. Recall that the SIT representation can be written as

$$V(\mathbb{F}) = \max_{F \in \mathbb{F}} (1 + K) U_0(F) - \sum_{s \in S} \max_{F \in \mathbb{F}} K \cdot U_s(F) \tag{A.14}$$

where

$$U_0(F) := \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega)u(f(\omega))$$

$$U_s(F) := \max_{g \in G} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega)u(g(\omega)) \text{ for each } s \in S$$

Observe that $U_0(F) = V(\{F\})$ for any menu $F \in \mathcal{M}$. Suppose $\{F\} \succsim \{G\}$ and $F \in \mathbb{F}$, we want to show that $\mathbb{F} \succsim \mathbb{F} \cup \{G\}$.

$\{F\} \succsim \{G\}$ implies that $U_0(F) \geq U_0(G)$. Together with $F \in \mathbb{F}$, this implies

$$\max_{H \in \mathbb{F}} (1+K)U_0(H) = \max_{H \in \mathbb{F} \cup \{G\}} (1+K)U_0(H).$$

That is, $V(\mathbb{F})$ and $V(\mathbb{F} \cup \{G\})$ have the same positive term in equation (A.14). On the other hand, for any $s \in S$, we have

$$\max_{H \in \mathbb{F}} K \cdot U_s(F) \leq \max_{H \in \mathbb{F} \cup \{G\}} K \cdot U_s(F).$$

Therefore, $V(\mathbb{F}) \geq V(\mathbb{F} \cup \{G\})$, indicating $\mathbb{F} \succsim \mathbb{F} \cup \{G\}$. □

Lemma A.17. *If $\succsim$ has a SIT representation, then $\succsim$ satisfies Axioms 1.6 (Interim Preference for Flexibility).*

Proof of Lemma A.17. We want to show that for any $\mathbb{F}$ and any F, G,

$$\mathbb{F} \cup \{F \cup G\} \succsim \mathbb{F} \cup \{F, G\}.$$

This is easier to see using the alternative expression for the IT representation:

$$\begin{aligned} V(\mathbb{F}) = & \max_{F \in \mathbb{F}} \left[(1+K) \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega)u(f(\omega)) \right] \\ & - K \sum_{s \in S} \max_{G \in \mathbb{F}} \max_{g \in G} \sum_{\omega \in \Omega} \pi(\omega)\sigma(s \mid \omega)u(g(\omega)) \end{aligned} \tag{A.15}$$

Note that the negative term can be equivalently written as

$$K \sum_{s\in S} \max_{g\in M(\mathbb{F})} \sum_{\omega\in\Omega} \pi(\omega)\sigma(s \mid \omega)u(g(\omega))$$

where $M(\mathbb{F}) = \{f \mid f \in F \text{ for some } F \in \mathbb{F}\}$ denotes the set of feasible acts as defined in the main text. It is clear that $M(\mathbb{F} \cup \{F \cup G\}) = M(\mathbb{F} \cup \{F, G\})$. Therefore, $V(\mathbb{F}\cup\{F\cup G\})$ and $V(\mathbb{F}\cup\{F, G\})$ have the same negative term in equation (A.15).

On the other hand, $V(\mathbb{F} \cup \{F \cup G\})$ has a weakly larger positive term because

$$\begin{aligned}&\sum_{s\in S} \max_{f\in F\cup G} \sum_{\omega\in\Omega} \pi(\omega)\sigma(s \mid \omega)u(f(\omega))\\ \geq\ &\max\left\{\sum_{s\in S} \max_{f\in F} \sum_{\omega\in\Omega} \pi(\omega)\sigma(s \mid \omega)u(f(\omega)), \sum_{s\in S} \max_{f\in G} \sum_{\omega\in\Omega} \pi(\omega)\sigma(s \mid \omega)u(f(\omega))\right\}\end{aligned}$$

Therefore, $V(\mathbb{F}\cup\{F\cup G\}) \geq V(\mathbb{F}\cup\{F, G\})$, indicating $\mathbb{F}\cup\{F\cup G\} \succsim \mathbb{F}\cup\{F, G\}$. □

Lemma A.18. *If $\succsim$ has a SIT representation, then $\succsim$ satisfies Axioms 1.7 and 1.8 (Nontriviality and Domination).*

Proof of Lemma A.18. For a lottery $\ell \in \Delta(X)$ and its corresponding constant act,

$$V(\{\{\ell\}\}) = \sum_{\omega\in\Omega} \pi(\omega)u(\ell) = u(\ell).$$

And Axiom 1.7 (Nontriviality) must be satisfied because we require u to be nonconstant.

For Axiom 1.8 (Domination), suppose f dominates g, then $f(\omega) \succsim g(\omega)$ for any $\omega \in \Omega$, which further implies that $u(f(\omega)) \geq u(g(\omega))$ for any $\omega \in \Omega$. Therefore,

$$V(\{\{f\}\}) = \sum_{\omega\in\Omega} \pi(\omega)u(f(\omega)) \geq \sum_{\omega\in\Omega} \pi(\omega)u(g(\omega)) = V(\{\{g\}\})$$

indicating $f \succsim g$. Moreover, $u(f(\omega)) \geq u(g(\omega))$ for any $\omega \in \Omega$ implies that

$$\sum_{\omega\in\Omega} \pi(\omega)\sigma(s \mid \omega)u(f(\omega)) \geq \sum_{\omega\in\Omega} \pi(\omega)\sigma(s \mid \omega)u(g(\omega)) \text{ for any } s \in S.$$

That is, g will never be chosen over f after any signal realization s. Therefore,

$$V(\{\{f,g\}\}) = \sum_{\omega\in\Omega} \pi(\omega)u(f(\omega)) = V(\{\{f\}\})$$

indicating $\{\{f,g\}\} \sim \{\{f\}\}$.

Similarly, if $\mathbb{F}$ dominates $\mathbb{G}$, then $\max_{F\in\mathbb{F}} U_0(F) = \max_{F\in\mathbb{F}\cup\mathbb{G}} U_0(F)$ and

$$\max_{F\in\mathbb{F}} U_s(F) = \max_{F\in\mathbb{F}\cup\mathbb{G}} U_s(F) \text{ for any } s \in S.$$

Therefore, $V(\mathbb{F}) = V(\mathbb{F}\cup\mathbb{G})$, indicating $\mathbb{F} \sim \mathbb{F}\cup\mathbb{G}$. □

This completes the proof for the necessity of Axioms 1.1-1.8 for a SIT representation.

A.2.2 Sufficiency of Axioms 1.1-1.8

As we have mentioned in the main text, the proof of Axioms 1.1-1.8 being sufficient for a SIT representation is more involved. Here we provide a roadmap before we dive into the details. We have characterized the partial regret (PR) representation in Appendix A.1. The PR representation looks very similar to the SIT representation, only with the caveat that it is in the framework of menus of menus of lotteries. In this proof, we will connect Axioms 1.1-1.8 we have in the main text with Axioms A.1-A.8 in the lottery framework in Appendix A.1 through two translations. The first translation will be from the preference over menus of menus of acts to menus of menus of "utility acts," and the second translation will be from the menus of menus of "utility acts" to menus of menus of lottery. Similar translation techniques are used in Dillenberger, Lleras, Sadowski, and Takeoka (2014).

Step 1: Translate the preference over directions to a preference over "utility directions."

Note that Axioms 1.1-1.3 imply their corresponding axioms (weak order, continuity and independence) over acts (i.e., direction containing only one singleton menu, like $\{\{f\}\}$). Axioms 1.7 (Nontriviality) and part 1 of Axiom 1.8 (Domination) imply the nontriviality axiom and the monotonicity axiom in the Anscombe-Aumann framework. Therefore, by result from the Anscombe-Aumann framework (for a detailed treatment, see Kreps (2018)), Axioms 1.1-1.3, 1.7 and 1.8 imply that there exists a unique probability measure $\bar{\pi} \in \Delta(\Omega)$ and a surjective affine utility index $u : \Delta(X) \to [0,1]$ such that the preference $\succsim$ restricting to acts can be represented by

$$V(\{\{f\}\}) := \sum_{\omega \in \Omega} \bar{\pi}(\omega) u\big(f(\omega)\big).$$

Now for any act $f \in \mathcal{F}$, the composite function $u \circ f : \Omega \to [0,1]$ specifies the utility associated with f in each state ω. We call this the *utility act* induced by f and also write $u(f)$ to denote the utility act induced by f.

Let $\mathcal{F}_u := u(\mathcal{F}) = \{u(f) \mid f \in \mathcal{F}\}$, that is, $\mathcal{F}_u$ is the set of utility acts induced by AA acts from $\mathcal{F}$. Since u is surjective, $\mathcal{F}_u = [0,1]^{|\Omega|}$.

Endow $\mathcal{F}_u$ with the Euclidean metric. Let $\mathcal{M}_u$ be the set of all non-empty compact subsets of $\mathcal{F}_u$, with typical elements F_u, G_u. We call these utility menus. Endow $\mathcal{M}_u$ with the Hausdorff metric. Let $\mathcal{D}_u$ be the set of all non-empty compact subsets of $\mathcal{M}_u$, with typical elements $\mathbb{F}_u, \mathbb{G}_u$. We call these utility directions. Endow $\mathcal{D}_u$ with the Hausdorff metric.

Naturally, we would believe that $\succsim$ as a binary relation over menus of menus of acts should induce a binary relation $\succsim_u$ over $\mathcal{D}_u$: If $\mathbb{F} \succsim \mathbb{G}$, then define $\mathbb{F}_u \succsim_u \mathbb{G}_u$ where $\mathbb{F}_u := u(\mathbb{F}) = \{u(F) \mid F \in \mathbb{F}\}$ and similar for $\mathbb{G}_u$. However, since u is generally not injective, for this definition to make sense, we need to guarantee that $u(\mathbb{F}) = u(\mathbb{G})$ implies $\mathbb{F} \sim \mathbb{G}$.

Lemma A.19. *If $\succsim$ satisfies Axioms 1.1-1.3, 1.7 and 1.8, then $u(\mathbb{F}) = u(\mathbb{G})$ implies $\mathbb{F} \sim \mathbb{G}$.*

Proof of Lemma A.19. We say two AA acts f and g are *indistinguishable* if

$$\{\{f(\omega)\}\} \sim \{\{g(\omega)\}\} \text{ for all } \omega \in \Omega.$$

Note that f and g are indistinguishable if and only if $u(f) = u(g)$. If this looks weird at first sight, recall that $u(f)$ and $u(g)$ are not scalar values of a function but are utility acts, that is, they are functions from Ω to $[0, 1]$.

We say two menus F and G are indistinguishable if for any $f \in F$ there exists $g \in G$ that is indistinguishable to f and for any $g' \in G$ there exists $f' \in F$ that is indistinguishable to g'. For a menu $F \in \mathcal{M}$, let $u(F) := \{u(f) \mid f \in F\}$. Note that F and G are indistinguishable if and only if $u(F) = u(G)$. (The only if part is easy. To see the if part, note that if $u(F) = u(G)$, then for any $f \in F$ there exists $g \in G$ such that $u(f) = u(g)$, which makes f and g indistinguishable. Similar arguments work for the other half of the arguments).

We say two directions $\mathbb{F}$ and $\mathbb{G}$ are indistinguishable if for any $F \in \mathbb{F}$ there exists $G \in \mathbb{G}$ that is indistinguishable to F and for any $G' \in \mathbb{G}$ there exists $F' \in \mathbb{F}$ that is indistinguishable to G'. For a direction $\mathbb{F} \in \mathcal{D}$, let $u(\mathbb{F}) := \{u(F) \mid F \in \mathbb{F}\}$. Then $\mathbb{F}$ and $\mathbb{G}$ are indistinguishable if and only if $u(\mathbb{F}) = u(\mathbb{G})$.

If $\mathbb{F}$ and $\mathbb{G}$ are indistinguishable, then $\mathbb{F}$ dominates $\mathbb{G}$ and $\mathbb{G}$ dominates $\mathbb{F}$, then by the second part of Axiom 1.8 (Domination),

$$\mathbb{F} \sim \mathbb{F} \cup \mathbb{G} \quad \text{and} \quad \mathbb{G} \sim \mathbb{F} \cup \mathbb{G}.$$

Therefore, $\mathbb{F} \sim \mathbb{G}$ by transitivty. □

We formally define a preference relation $\succsim_u$ over $\mathcal{D}_u$ by

$$\mathbb{F}_u \succsim_u \mathbb{G}_u \text{ if and only if } \mathbb{F} \succsim \mathbb{G} \text{ where } \mathbb{F} \in u^{-1}(\mathbb{F}_u) \text{ and } \mathbb{G} \in u^{-1}(\mathbb{G}_u). \tag{A.16}$$

This is a valid definition by the result of Lemma A.19. We then show that if $\succsim$ satisfies Axioms 1.1-1.8, then $\succsim_u$ satisfies the suitably adapted versions of Axioms 1.1-1.8. The formal description of the adapted axioms are as below.

Axiom A.2.1. *$\succsim_u$ is complete and transitive.*

Axiom A.2.2. *For any $\mathbb{F}_u$, the sets $\{\mathbb{G}_u : \mathbb{G}_u \succsim_u \mathbb{F}_u\}$, $\{\mathbb{G}_u : \mathbb{F}_u \succsim_u \mathbb{G}_u\}$ are closed.*

Axiom A.2.3. *For any $\mathbb{F}_u, \mathbb{G}_u, \mathbb{H}_u$ and any $\alpha \in (0,1)$,*

$$\mathbb{F}_u \succsim_u \mathbb{G}_u \iff \alpha\mathbb{F}_u + (1-\alpha)\mathbb{H}_u \succsim_u \alpha\mathbb{G}_u + (1-\alpha)\mathbb{H}_u.$$

We can define the notion of a critical direction and a critical menu within a direction with respect to $\succsim_u$ in a similar way that is defined in the main text.

Axiom A.2.4. *There exists a natural number N such that*

- *For every $\mathbb{F}_u \in \mathcal{D}_u$, there exists $\mathbb{G}_u$ with $|\mathbb{G}_u| < N$ such that $\mathbb{G}_u$ is critical for $\mathbb{F}_u$;*
- *For every $\mathbb{F}_u \in \mathcal{D}_u$ and every $F_u \in \mathbb{F}_u$, there exists G_u with $|G_u| < N$ such that G_u is critical for F_u in $\mathbb{F}_u$.*

Axiom A.2.5. *If $\{F_u\} \succsim_u \{G_u\}$ and $F_u \in \mathbb{F}_u$, then $\mathbb{F}_u \succsim \mathbb{F}_u \cup \{G_u\}$.*

Axiom A.2.6. *For any $\mathbb{F}_u$ and any F_u, G_u, $\mathbb{F}_u \cup \{F_u \cup G_u\} \succsim_u \mathbb{F}_u \cup \{F_u, G_u\}$.*

Axiom A.2.7. *There exist $\mathbb{F}_u, \mathbb{G}_u \in \mathcal{D}_u$ such that $\mathbb{F}_u \supseteq \mathbb{G}_u$ and $\mathbb{F}_u \succ_u \mathbb{G}_u$.*

We can define the notion of domination with respect to utility acts, utility menus and utility directions in a similar way those are defined in the main text.

Axiom A.2.8.

- *If f_u dominates g_u, then $f_u \succsim_u g_u$ and $\{\{f_u, g_u\} \sim_u \{\{f_u\}\}$;*
- *If $\mathbb{F}_u$ dominates $\mathbb{G}_u$, then $\mathbb{F}_u \sim_u \mathbb{F}_u \cup \mathbb{G}_u$.*

Lemma A.20. *If a binary relation $\succsim$ over directions satisfies Axioms 1.1-1.8, then the induced binary relation $\succsim_u$ over utility directions satisfies Axioms A.2.1-A.2.8.*

Proof. See Appendix A.4.1. □

This completes Step 1 in our translation.

Step 2: Translate the preference over utility directions to a preference over menus of menus of lotteries.

Recall that $\mathcal{F}_u$ is the collection of all utility acts and we can identify $\mathcal{F}_u$ with the set of all n-dimensional vectors where each entry is in $[0,1]$. That is, $\mathcal{F} = [0,1]^n$, where $n = |\Omega|$ is the cardinality of the state space Ω.

For convenience, let $\Omega = \{\omega_1, \ldots, \omega_n\}$, and introduce an artificial state ω_0 and a new space $\mathcal{F}'$ defined by

$$\mathcal{F}' := \left\{ f' \in [0,n] \times [0,1]^n \;\middle|\; \sum_{i=0}^{n} f'(\omega_i) = n \right\}. \tag{A.17}$$

For notational simplicity, we suppress the subscript u but one should keep in mind that these are still interepreted as utility acts.

Endow $\mathcal{F}'$ with the standard Euclidean metric. Consider $r : \mathcal{F}' \to \mathcal{F}_u$ where $r(f')$ is the vector in $\mathcal{F}$ that agrees with the last n components of f', that is,

$$r(f') \in \mathcal{F}_u \text{ with } [r(f')](\omega_i) = f'(\omega_i),\ \forall i \in \{1, 2, \ldots, n\}.$$

It is easy to verify that r is a homeomorphism between $\mathcal{F}'$ and $\mathcal{F}_u$. Let r^{-1} denote its inverse.

Let $\mathcal{M}'$ denote the set of nonempty compact subsets of $\mathcal{F}'$, with typical elements F', G'. Endow $\mathcal{M}'$ with the Hausdorff metric. Let $\mathcal{D}'$ denote the set of nonempty compact subsets of $\mathcal{M}'$, with typical elements $\mathbb{F}', \mathbb{G}'$. Endow $\mathcal{D}'$ with the Hausdorff metric.

We slightly abuse notation and let r also denote the homeomorphism from $\mathcal{M}'$ to $\mathcal{M}_u$ and $\mathcal{D}'$ to $\mathcal{D}_u$, with $r(F') = \{r(f') \mid f' \in F'\} \in \mathcal{M}_u$ and $r(\mathbb{F}') = \{r(F') \mid F' \in \mathbb{F}'\} \in \mathcal{D}_u$.

With this construction, we can define a preference $\succsim_*$ over $\mathcal{D}'$ by

$$\mathbb{F}' \succsim_* \mathbb{G}' \text{ if } r(\mathbb{F}') \succsim_u r(\mathbb{G}').$$

To move on, we introduce yet another new space by

$$\mathcal{F}'' := \left\{ f'' \in [0,n]^{n+1} \;\middle|\; \sum_{i=0}^{n} f''(\omega_i) = n \right\}.$$

Endow $\mathcal{F}''$ with the standard Euclidean metric. Let $\mathcal{M}''$ denote the set of nonempty compact subsets of $\mathcal{F}''$, with typical elements F'', G''. Endow $\mathcal{M}''$ with the Hausdorff metric. Let $\mathcal{D}''$ denote the set of nonempty compact subsets of $\mathcal{M}''$, with typical elements $\mathbb{F}'', \mathbb{G}''$. Endow $\mathcal{D}''$ with the Hausdorff metric.

Fix a menu F^{n+1} defined by

$$F^{n+1} := \left\{ \left(\frac{n}{n+1}, \ldots, \frac{n}{n+1} \right) \right\}.$$

Then $F^{n+1} \in \mathcal{M}'$. Also observe that for any $F'' \in \mathcal{M}''$ and $\varepsilon \leq \frac{1}{n^2}$,

$$\varepsilon F'' + (1-\varepsilon)F^{n+1} = \left\{ \varepsilon f'' + (1-\varepsilon) \left(\frac{n}{n+1}, \ldots, \frac{n}{n+1} \right) \;\middle|\; f'' \in F'' \right\} \in \mathcal{M}'.$$

Finally, define a relation $\succsim_{**}$ on $\mathcal{D}''$ by: $\mathbb{F}'' \succsim_{**} \mathbb{G}''$ if

$$\varepsilon \mathbb{F}'' + (1-\varepsilon)\{F^{n+1}\} \succsim_{*} \varepsilon \mathbb{G}'' + (1-\varepsilon)\{F^{n+1}\} \text{ for all } \varepsilon < \frac{1}{n^2}.$$

We then show that if $\succsim_u$ satisfies Axioms A.2.1-A.2.8, then $\succsim_*$ satisfies the suitably adapted versions of these axioms. These adapted axioms are listed as Axioms A.2.9-A.2.16.

Axiom A.2.9. $\succsim_*$ *is complete and transitive.*

Axiom A.2.10. *For any* $\mathbb{F}'$*, the sets* $\{\mathbb{G}' : \mathbb{G}' \succsim_* \mathbb{F}'\}$*,* $\{\mathbb{G}' : \mathbb{F}' \succsim_* \mathbb{G}'\}$ *are closed.*

Axiom A.2.11. *For any* $\mathbb{F}', \mathbb{G}', \mathbb{H}'$ *and any* $\alpha \in (0,1)$*,*

$$\mathbb{F}' \succsim_* \mathbb{G}' \iff \alpha\mathbb{F}' + (1-\alpha)\mathbb{H}' \succsim_* \alpha\mathbb{G}' + (1-\alpha)\mathbb{H}'.$$

We can define the notion of a critical direction and a critical menu within a direction with respect to $\succsim_*$ in a similar way as before.

Axiom A.2.12. *There exists a natural number N such that*

- *For every $\mathbb{F}' \in \mathcal{D}'$, there exists $\mathbb{G}'$ with $|\mathbb{G}'| < N$ such that $\mathbb{G}'$ is critical for $\mathbb{F}'$;*
- *For every $\mathbb{F}' \in \mathcal{D}'$ and every $F' \in \mathbb{F}'$, there exists G' with $|G'| < N$ such that G' is critical for F' in $\mathbb{F}'$.*

Axiom A.2.13. *If $\{F'\} \succsim_* \{G'\}$ and $F' \in \mathbb{F}'$, then $\mathbb{F}' \succsim_* \mathbb{F}' \cup \{G'\}$.*

Axiom A.2.14. *For any $\mathbb{F}' \in \mathcal{D}'$ and any $F', G' \in \mathcal{M}'$, $\mathbb{F}' \cup \{F' \cup G'\} \succsim_* \mathbb{F}' \cup \{F', G'\}$.*

Axiom A.2.15. *There exist $\mathbb{F}', \mathbb{G}' \in \mathcal{D}'$ such that $\mathbb{F}' \supseteq \mathbb{G}'$ and $\mathbb{F}' \succ_* \mathbb{G}'$.*

Axiom A.2.16. *If $G' \subseteq F'$ and $F' \in \mathbb{F}'$, then $\mathbb{F}' \cup \{G'\} \succsim_* \mathbb{F}'$.*

Note that the expression of Axiom A.2.16 does not explicitly appear in Axioms 1.1-1.8, it is implied from part 2 of Axiom A.2.8 and closely related to Axiom A.7.

Lemma A.21. *If a binary relation $\succsim_u$ over utility directions satisfies Axioms A.2.1-A.2.8, then the induced binary relation $\succsim_*$ over $\mathcal{D}'$ satisfies Axioms A.2.9-A.2.16.*

Proof. See Appendix A.4.2. □

Step 3: Verify that $\succsim_{**}$ is the unique extension of $\succsim_*$ to $\mathcal{D}''$ satisfying adapted versions of Axioms A.2.9-A.2.16 if $\succsim_*$ satisfies Axioms A.2.9-A.2.16.

The construction of $\succsim_*$ and $\succsim_{**}$ in Step 2 is the same as that in Dillenberger, Lleras, Sadowski, and Takeoka (2014).

We proceed by first listing the adapted versions of Axioms A.2.9-A.2.16. These are listed as Axioms A.2.17-A.2.24.

Axiom A.2.17. *$\succsim_{**}$ is complete and transitive.*

Axiom A.2.18. *For any $\mathbb{F}''$, the sets $\{\mathbb{G}'' : \mathbb{G}'' \succsim_{**} \mathbb{F}''\}$, $\{\mathbb{G}'' : \mathbb{F}'' \succsim_{**} \mathbb{G}''\}$ are closed.*

Axiom A.2.19. *For any* $\mathbb{F}'', \mathbb{G}'', \mathbb{H}''$ *and any* $\alpha \in (0,1)$,

$$\mathbb{F}'' \succsim_{**} \mathbb{G}'' \iff \alpha\mathbb{F}'' + (1-\alpha)\mathbb{H}'' \succsim_{**} \alpha\mathbb{G}'' + (1-\alpha)\mathbb{H}''.$$

We can define the notion of a critical direction and a critical menu within a direction with respect to $\succsim_{**}$ in a similar way as before.

Axiom A.2.20. *There exists a natural number* N *such that*

- *For every* $\mathbb{F}'' \in \mathcal{D}''$*, there exists* $\mathbb{G}''$ *with* $|\mathbb{G}''| < N$ *such that* $\mathbb{G}''$ *is critical for* $\mathbb{F}''$*;*
- *For every* $\mathbb{F}'' \in \mathcal{D}''$ *and every* $F'' \in \mathbb{F}''$*, there exists* G'' *with* $|G''| < N$ *such that* G'' *is critical for* F'' *in* $\mathbb{F}''$*.*

Axiom A.2.21. *If* $\{F''\} \succsim_{**} \{G''\}$ *and* $F'' \in \mathbb{F}''$*, then* $\mathbb{F}'' \succsim_{**} \mathbb{F}'' \cup \{G''\}$*.*

Axiom A.2.22. *For any* $\mathbb{F}'' \in \mathcal{D}''$ *and any* $F'', G'' \in \mathcal{M}''$*,* $\mathbb{F}'' \cup \{F'' \cup G''\} \succsim_{**} \mathbb{F}'' \cup \{F'', G''\}$*.*

Axiom A.2.23. *There exist* $\mathbb{F}'', \mathbb{G}'' \in \mathcal{D}''$ *such that* $\mathbb{F}'' \supseteq \mathbb{G}''$ *and* $\mathbb{F}'' \succ_{**} \mathbb{G}''$*.*

Axiom A.2.24. *If* $G'' \subseteq F''$ *and* $F'' \in \mathbb{F}''$*, then* $\mathbb{F}'' \cup \{G''\} \succsim_{**} \mathbb{F}''$*.*

Lemma A.22. *If a binary relation* $\succsim_*$ *over* $\mathcal{D}'$ *satisfies Axioms A.2.9-A.2.16, then the induced binary relation* $\succsim_{**}$ *over* $\mathcal{D}''$ *satisfies Axioms A.2.17-A.2.24.*

Proof. See Appendix A.4.3. □

Step 4: Apply Theorem A.2 in Appendix A and translate the resulting PR representation to a SIT representation.

We are very close to a SIT representation. We proceed by first rescaling every element of $\mathcal{F}''$ with factor $\frac{1}{n}$. The rescaling will give us a unit simplex and make the corresponding domain $\mathcal{D}''$ formally equivalent to the choice domain in Appendix A.1.

Through steps 2 and 3, we have verified that if a binary relation $\succsim$ over the direction of acts satisfies Axioms 1.1-1.8, then the binary relation $\succsim_{**}$ defined over $\mathcal{D}''$ in step 2 satisfies Axioms A.2.17-A.2.24, and it is uniquely derived from the original binary relation $\succsim$ over directions of acts. Therefore, we can apply Theorem A.2 to conclude that there exists a finitely-supported probability measure $\widehat{\mu}$ over $\widehat{\mathcal{U}}$ (the set of doubly normalized expected utilities over $\Delta(\widehat{\Omega})$ where $\widehat{\Omega} = \Omega \cup \{\omega_0\}$) and a scalar $\widehat{K} \geq 0$ such that $\succsim_{**}$ can be represented by

$$\widehat{W}(\mathbb{F}'') = \max_{F'' \in \mathbb{F}''} \left[(1+\widehat{K}) \sum_{\widehat{u} \in \text{supp}(\widehat{\mu})} \widehat{\mu}(\widehat{u}) \max_{f'' \in F''} \widehat{u}(f'') \right] - \widehat{K} \sum_{\widehat{u} \in \text{supp}(\widehat{\mu})} \widehat{\mu}(\widehat{u}) \max_{F'' \in \mathbb{F}''} \max_{f'' \in F''} \widehat{u}(f'') \tag{A.18}$$

where $\widehat{u}(f'')$ is the expected utility for lottery f'' under taste $\widehat{u}$. That is,

$$\widehat{u}(f'') = \sum_{\widehat{\omega} \in \widehat{\Omega}} \widehat{u}(\widehat{\omega}) f''(\widehat{\omega}) \text{ for any } \widehat{u} \in \widehat{\mathcal{U}} \text{ and } f'' \in \mathcal{F}''.$$

Moreover, $\widehat{\mu}$ is uniquely identified (Lemma A.9), and $\widehat{K}$ is uniquely identified when $|\text{supp}(\widehat{\mu})| > 1$ (Theorem A.13). Note that $\widehat{W}$ also represents $\succsim_*$ when restricting to $\mathcal{D}'$.

To get to the SIT representation, we aim for a representation for $\succsim$ of the form

$$V(\mathbb{F}) = \max_{F \in \mathbb{F}} \left[(1+K) \sum_{\pi \in \text{supp}(\nu)} \nu(\pi) \max_{f \in F} u_\pi(f_u) \right] - K \sum_{\pi \in \text{supp}(\nu)} \nu(\pi) \max_{F \in \mathbb{F}} \max_{f \in F} u_\pi(f_u) \tag{A.19}$$

where ν is a finitely-supported probability measure over $\Delta(\Omega)$ representing a distribution over posteriors induced by a prior and an information structure, and

$$u_\pi(f_u) := \sum_{\omega \in \Omega} \pi(\omega) f_u(\omega) \text{ for all } f_u \in \mathcal{F}_u \text{ and } \pi \in \Delta(\Omega).$$

We now explore the additional constraint imposed on $\widehat{W}$ by Axiom A.2.8.

Lemma A.23. *Suppose $\succsim_{**}$ has a representation as defined in equation* (A.18) *and $\succsim_u$ satisfies Axiom A.2.8, then $\widehat{u}(\omega) \geq \widehat{u}(\omega_0)$ for any $\widehat{u} \in \text{supp}(\widehat{\mu})$ and any $\omega \in \Omega$.*

Proof of Lemma A.23. Suppose by contradiction that $\succsim_{**}$ can be represented as in equation (A.18), $\succsim_u$ satisfies Axiom A.2.8, but there exists $\widehat{u}_* \in \text{supp}(\widehat{\mu})$ and $\omega_* \in \Omega$ such that $\widehat{u}_*(\omega_0) > \widehat{u}_*(\omega_*)$. We want to derive a contradiction. Let

$$f' := (n - \varepsilon, 0, \ldots, 0, \varepsilon, 0, \ldots, 0) \tag{A.20}$$

where $n - \varepsilon$ is assigned to state ω_0, ε is assigned to state ω_* and 0 is assigned to any other state. We can find ε small enough so that $f' \in \mathcal{F}'$. Let

$$g' := (n, 0, \ldots, 0) \tag{A.21}$$

where n is assigned to state ω_0 and 0 is assigned to any state $\omega \in \Omega$. Therefore, $r(f')$ dominates $r(g')$ (recall that the notion of domination is defined over $\mathcal{F}_u$ and thus only concerns the last n coordinates).

Thus, part 1 of Axiom A.2.8 dictates that $\{\{r(f')\}\} \sim_u \{\{r(f'), r(g')\}\}$, and by the definition of $\succsim_*$, we must have $\{\{f'\}\} \sim_* \{\{f', g'\}\}$.

Note that for any $\widehat{u} \in \text{supp}(\widehat{\mu})$,

$$\widehat{u}(f') = (n - \varepsilon)\widehat{u}(\omega_0) + \varepsilon \widehat{u}(\omega_*)$$
$$\widehat{u}(g') = n\widehat{u}(\omega_0)$$

In particular,

$$\widehat{u}_*(f') - \widehat{u}_*(g') = \varepsilon\Big(\widehat{u}_*(\omega_*) - \widehat{u}_*(\omega_0)\Big) < 0.$$

Therefore,

$$\begin{aligned}
\widehat{W}(\{\{f', g'\}\}) &= \sum_{\widehat{u} \in \text{supp}(\widehat{\mu})} \widehat{\mu}(\widehat{u}) \Big[\widehat{u}(f') \vee \widehat{u}(g')\Big] \\
&= \sum_{\widehat{u} \neq \widehat{u}_*} \widehat{\mu}(\widehat{u}) \Big[\widehat{u}(f') \vee \widehat{u}(g')\Big] + \widehat{\mu}(\widehat{u}_*)\widehat{u}_*(g') \\
&> \sum_{\widehat{u} \neq \widehat{u}_*} \widehat{\mu}(\widehat{u})\widehat{u}(f') + \widehat{\mu}(\widehat{u}_*)\widehat{u}_*(f') \\
&= \widehat{W}(\{\{f'\}\})
\end{aligned}$$

contradicting $\{\{f', g'\}\} \sim_* \{\{f'\}\}$. This completes the proof of Lemma A.23. $\square$

Given our construction of $\widehat{W}$, there are two normalizations we can do that will allow us to replace the finitely-supported probability measure $\widehat{\mu}$ on $\widehat{\mathcal{U}}$ with a unique finitely-supported probability measure ν on $\Delta(\Omega)$.

For all $\omega \in \Omega$ and for all $\widehat{u}$, define

$$\xi(\widehat{u}(\omega)) := \widehat{u}(\omega) - \widehat{u}(\omega_0).$$

Since $\sum_{i=0}^{n} f'(\omega_i) = n$ for all $f' \in \mathcal{F}'$ and ξ simply adds a constant to every $\widehat{u}$,

$$\arg\max_{f'' \in F''} \left(\sum_{i=0}^{n} f''(\omega_i) \xi(\widehat{u}(\omega_i)) \right) = \arg\max_{f'' \in F''} \left(\sum_{i=1}^{n} f''(\omega_i) \widehat{u}(\omega_i) \right)$$

for all $F'' \in \mathcal{M}''$ and all $\widehat{u} \in \operatorname{supp}(\widehat{\mu})$. Furthermore, by Lemma A.23, $\xi(\widehat{u}(\omega)) \geq 0$ for all $\omega \in \Omega$ and all $\widehat{u} \in \operatorname{supp}(\widehat{\mu})$.

Therefore, we could transform $\xi \circ \widehat{u} : \Omega \to \mathbb{R}$ into a probability measure $\pi_{\widehat{u}}$. This can be done by letting $\pi_{\widehat{u}} \in \Delta(\Omega)$ be defined by

$$\pi_{\widehat{u}}(\omega) := \frac{\xi(\widehat{u}(\omega))}{\sum_{\omega' \in \Omega} \xi(\widehat{u}(\omega'))} \text{ for each } \omega \in \Omega.$$

Finally, we can get the SIT representation by adjusting the weights on each $\pi_{\widehat{u}}$ by defining $\nu \in \Delta_0(\Delta(\Omega))$ by

$$\nu(\pi_{\widehat{u}}) := \frac{\sum_{\omega \in \Omega} \xi(\widehat{u}(\omega))}{\sum_{\widehat{u}' \in \operatorname{supp}(\hat{\mu})} \sum_{\omega \in \Omega} \xi(\widehat{u}'(\omega))} \widehat{\mu}(\widehat{u}).$$

This will give us the representation we aim for.

Lastly, we can let the set of signal realizations S equal to $\{1, 2, \ldots, m\}$ where $m = |\operatorname{supp}(\nu)|$ is the size of the support for ν, and the conditional probability distributions $\sigma(s \mid \omega)$ can be uniquely determined.

A.3 Other Omitted Proofs in Chapter 1

A.3.1 Proof of Theorem 1.2

Proof. The uniqueness of the SIT representation follows from the uniqueness of the PR representation. Note that through the proof of Theorem 1.1, we have established that if $\succsim$ has a SIT representation with parameters (π, u, K, σ), then the translated preference $\succsim_{**}$ has a PR representation with parameters (ν, K) where ν is a distribution over posteriors induced by π and σ and K is a non-negative regret intensity level.

Therefore, if $(\pi_0, u_0, K_0, \sigma_0)$ and (π, u, K, σ) both represent $\succsim$ through a SIT representation, then (ν_0, K_0) and (ν, K) both represent the induced preference $\succsim_{**}$ through a PR representation. Then we can apply Lemma 13 and Theorem A.13 to conclude that $\nu_0 = \nu$ and $K_0 = K$ whenever $|\text{supp}(\nu)| > 1$. This completes the proof. $\square$

A.3.2 Proof of Lemma 1.4

Proof.

Only if. Suppose $\succsim$ has an aligned IT representation, then $\succsim$ must be complete and transitive. Moreover, each $\succsim^\sigma$ has a SIT representation with parameters $(\pi, u, i^\sigma, K^\sigma)$, then by Theorem 1.1, $\succsim^\sigma$ satisfies Axioms 1.2-1.8 for each $\sigma \in \mathcal{I}$. We now verify that $\succsim$ satisfies Stable Preference over Acts and Act Independence.

First note that for any act $f \in \mathcal{F}$ and any $\sigma \in \mathcal{I}$, the value of f under W is independent of σ since

$$W(f, \sigma) = \sum_{\omega \in \Omega} \pi(\omega) u(f(\omega)).$$

Therefore, Stable Preference over Acts must be satisfied.

To show that $\succsim$ satisfies Act Independence, note that W is affine in its first argument. That is, for any $\mathbb{F}, \mathbb{G} \in \mathcal{D}$ and any fixed information structure $\sigma \in \mathcal{I}$,

$$W(\alpha\mathbb{F} + (1-\alpha)\mathbb{G}, \sigma) = \alpha W(\mathbb{F}, \sigma) + (1-\alpha)W(\mathbb{G}, \sigma).$$

Therefore, for any $\mathbb{F}, \mathbb{G} \in \mathcal{D}$, $\sigma, \sigma' \in \mathcal{I}$ and $\alpha \in (0,1)$,

$$\begin{aligned}
& (\alpha\mathbb{F} + (1-\alpha)h, \sigma) \succsim (\alpha\mathbb{G} + (1-\alpha)h, \sigma') \\
\iff & W(\alpha\mathbb{F} + (1-\alpha)h, \sigma) \geq W(\alpha\mathbb{G} + (1-\alpha)h, \sigma') \\
\iff & \alpha W(\mathbb{F}, \sigma) + (1-\alpha)W(h, \sigma) \geq \alpha W(\mathbb{G}, \sigma') + (1-\alpha)W(h, \sigma') \\
\iff & W(\mathbb{F}, \sigma) \geq W(\mathbb{G}, \sigma') \\
\iff & (\mathbb{F}, \sigma) \succsim (\mathbb{G}, \sigma')
\end{aligned}$$

where the penultimate equivalence follows from the above result that $W(h, \sigma) = W(h, \sigma')$.

This completes the proof of the "only if" part.

If. Suppose $\succsim$ satisfies Weak Order, Stable Preference for Acts and Act Independence, and $\succsim^\sigma$ satisfies Axioms 1.2-1.8 for each $\sigma \in \mathcal{I}$, we want to show that $\succsim$ has an aligned informational tradeoff representation.

As mentioned in the main text, $\succsim$ being complete and transitive implies that $\succsim^\sigma$ is complete and transitive for each $\sigma \in \mathcal{I}$. Therefore, we can apply Theorem 1.1 to conclude that each conditional $\succsim^\sigma$ has a SIT representation V^σ with parameters $(\pi^\sigma, u^\sigma, K^\sigma, i^\sigma)$.

Lemma A.24. *If $\succsim$ satisfies Stable Preference over Acts, then there exists $\pi \in \Delta(\Omega)$ and $u : \Delta(X) \to \mathbb{R}$ such that $\pi^\sigma = \pi$ and $u^\sigma = \alpha^\sigma u + b^\sigma$ for some $\alpha^\sigma > 0$ and $b^\sigma \in \mathbb{R}$ for all $\sigma \in \mathcal{I}$.*

Proof of Lemma A.24. Note that when restricting to acts,

$$V^\sigma(f) = \sum_{\omega \in \Omega} \pi^\sigma(\omega) u^\sigma(f(\omega)) \text{ for any act } f \in \mathcal{F}.$$

That is, each V^σ reduces to a subjective expected utility representation with parameters (π^σ, u^σ) when restricting to acts. Suppose $\succsim$ satisfies Stable Preference over Acts, then $f \succsim^\sigma g$ if and only if $f \succsim^{\sigma'} g$ for any acts f, g and any information structures σ, σ'. Therefore, $\succsim^\sigma$ and $\succsim^{\sigma'}$ induce the same preference over acts. Then the result follows directly from the uniqueness result of the Anscombe-Aumann framework (see Kreps (2018) for one treatment). □

We are not quite done yet as we have not established that $\succsim$ actually has a utility representation. We can take care of this using the other axiom, Act Independence. Let $\widetilde{V}^\sigma$ be a SIT representation with parameters $(\pi, u, K^\sigma, i^\sigma)$ where π, u are as characterized in Lemma A.24. Then $\widetilde{V}^\sigma$ represents $\succsim^\sigma$ for each $\sigma \in \mathcal{I}$.

Consider a function $W : \mathcal{D} \times \mathcal{I} \to \mathbb{R}$ defined by

$$W(\mathbb{F}, \sigma) := \widetilde{V}^\sigma(\mathbb{F}).$$

Suppose $\succsim$ also satisfies Act Independence, we will now show that $\succsim$ can be represented by W. That is, we want to show that $(\mathbb{F}, \sigma) \succsim (\mathbb{G}, \sigma')$ if and only if $W(\mathbb{F}, \sigma) \geq W(\mathbb{G}, \sigma')$. Fix two pairs $(\mathbb{F}, \sigma)$ and $(\mathbb{G}, \sigma')$.

Lemma A.25. *There exist acts f, g, h and $\alpha \in (0, 1]$ such that*

$$(\alpha\mathbb{F} + (1 - \alpha)h, \sigma) \sim (f, \sigma)$$
$$(\alpha\mathbb{G} + (1 - \alpha)h, \sigma') \sim (g, \sigma')$$

Proof of Lemma A.25. It is without loss to normalize the taste function $u : \Delta(X) \to \mathbb{R}$ to have range $[0, 1]$ (one way to do this is to assign utility 0 to the worst outcome in X and utility 1 to the best outcome in X and this is always doable since X is finite). Moreover, $u : \Delta(X) \to [0, 1]$ is surjective, so for any $a \in [0, 1]$, there exists an act $f \in \mathcal{F}$ such that $\widetilde{V}^\sigma(f) = a$ for all $\sigma \in \mathcal{I}$. Using this, we can show that the value of any menu $F \in \mathcal{M}$ under any information structure is in $[0, 1]$, that is, $0 \leq \widetilde{V}^\sigma(\mathbb{F}) \leq 1$

for any $F \in \mathcal{M}$ and $\sigma \in \mathcal{I}$. Finally, this indicates that $\widetilde{V}^\sigma(\mathbb{F}) \leq \widetilde{V}^\sigma(M(\mathbb{F})) \leq 1$ for any direction $\mathbb{F}$ and information structure $\sigma \in \mathcal{I}$. Therefore, for any act h and any $\sigma \in \mathcal{I}$,

$$\widetilde{V}^\sigma\Big(\alpha\mathbb{F} + (1-\alpha)h\Big) = \alpha\widetilde{V}^\sigma(\mathbb{F}) + (1-\alpha)\widetilde{V}^\sigma(h)$$

$$\widetilde{V}^{\sigma'}\Big(\alpha\mathbb{G} + (1-\alpha)h\Big) = \alpha\widetilde{V}^{\sigma'}(\mathbb{G}) + (1-\alpha)\widetilde{V}^{\sigma'}(h)$$

By our arguments above, there exists h such that $\widetilde{V}^\sigma(h) = \widetilde{V}^{\sigma'}(h) = 1$. Therefore for such an h, $\widetilde{V}^\sigma(h) \geq \widetilde{V}^\sigma(\mathbb{F})$ and $\widetilde{V}^{\sigma'}(h) \geq \widetilde{V}^{\sigma'}(\mathbb{G})$, which further implies that we can find $\alpha^* \in (0,1]$ such that

$$\widetilde{V}^\sigma\Big(\alpha^*\mathbb{F} + (1-\alpha^*)h\Big) = \alpha^*\widetilde{V}^\sigma(\mathbb{F}) + (1-\alpha^*)\widetilde{V}^\sigma(h) \geq 0$$

$$\widetilde{V}^{\sigma'}\Big(\alpha^*\mathbb{G} + (1-\alpha^*)h\Big) = \alpha^*\widetilde{V}^{\sigma'}(\mathbb{G}) + (1-\alpha^*)\widetilde{V}^{\sigma'}(h) \geq 0$$

By the surjectivity of u, this guarantees the existence of some acts $f, g \in \mathcal{F}$ such that

$$(\alpha^*\mathbb{F} + (1-\alpha^*)h, \sigma) \sim (f, \sigma)$$

$$(\alpha^*\mathbb{G} + (1-\alpha^*)h, \sigma') \sim (g, \sigma')$$

This completes the proof. □

Fix two pairs $(\mathbb{F}, \sigma)$ and $(\mathbb{G}, \sigma')$. Let f, g, h and $\alpha \in (0,1]$ be as characterized from Lemma A.25. Then

$$\begin{aligned}
(\mathbb{F}, \sigma) \succsim (\mathbb{G}, \sigma') &\iff (\alpha\mathbb{F} + (1-\alpha)h, \sigma) \succsim (\alpha\mathbb{G} + (1-\alpha)h, \sigma') \\
&\iff (f, \sigma) \succsim (g, \sigma') \\
&\iff (f, \sigma) \succsim (g, \sigma) \\
&\iff \sum_{\omega \in \Omega} \pi(\omega)u(f(\omega)) \geq \sum_{\omega \in \Omega} \pi(\omega)u(g(\omega)) \\
&\iff W(\alpha\mathbb{F} + (1-\alpha)h, \sigma) \geq W(\alpha\mathbb{G} + (1-\alpha)h, \sigma') \\
&\iff W(\mathbb{F}, \sigma) \geq W(\mathbb{G}, \sigma')
\end{aligned}$$

The first equivalence holds since $\succsim$ satisfies Act Independence. The second equivalence holds by the construction of f, g, h and α. The third equivalence follows from $\succsim$ satisfying Stable Preference for Acts. The fourth equivalence holds since $\widetilde{V}^\sigma$ represents $\succsim^\sigma$. The fifth equivalence holds since $W(\mathbb{F}, \sigma) = \widetilde{V}^\sigma(\mathbb{F})$ so $W(\cdot, \sigma)$ can represent $\succsim^\sigma$ and $(\alpha\mathbb{F} + (1-\alpha)h, \sigma) \sim (f, \sigma)$ by construction. The last equivalence holds since W is affine in the directions.

This completes the proof of the sufficiency of the axioms since W is by definition an aligned informational tradeoff representation. □

A.3.3 Proof of Lemma 1.5

Proof. If. If i^o induces a degenerate distribution over posteriors, $\succsim^o$ is represented by

$$V^o(\mathbb{F}) = \max_{F \in \mathbb{F}} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega) u(f(\omega)). \tag{A.22}$$

If $V^o(\{F\}) \geq V^o(\{G\})$, then $V^o(\{F \cup G\}) = V^o(\{F\})$, so Strategic Rationality when No Information (SRNI) is satisfied.

Only if. Suppose i^o induces a non-degenerate distribution over posteriors, that is, the support of the induced distribution over posteriors contains at least two different elements, μ_1 and μ_2. Then by a standard result, there exist two acts f and g such that f yields a strictly higher expected utility than g given posterior μ_1 and g yields a strictly higher expected utility than f given posterior μ_2. Therefore, $V^o(\{\{f, g\}\}) > \max\{V^o(\{\{f\}\}), V^o(\{\{g\}\})\}$, violating SRNI. Taking the contrapositive completes the proof. □

A.3.4 Proof of Lemma 1.6

Proof. Only if. Suppose $\succsim$ has a regret-varying IT representation, then $\succsim$ has an aligned informational tradeoff representation with $i^\sigma = \sigma$ for each $\sigma \in \mathcal{I}$. Therefore, we can apply Lemma 1.4 to conclude that $\succsim$ satisfies Weak Order and Act Independence, and $\succsim^\sigma$ satisfies Axioms 1.2-1.8 for each $\sigma \in \mathcal{I}$. Further more, $\succsim^o$ can be represented by (π, u, K^o, o), therefore, we can apply Lemma 1.5 to conclude that $\succsim$ satisfies SRNI.

We now check that $\succsim$ must satisfy Reduction. Note that the regret terms are zero for singleton directions $\{F\}$ and $\{F_\sigma\}$. Moreover,

$$\begin{aligned}
W(\{F_\sigma\}, o) &= \max_{f \in F_\sigma} \sum_{\omega \in \Omega} \pi(\omega) u(f(\omega)) \\
&= \max_{\gamma \in F^S} \sum_{\omega \in \Omega} \pi(\omega) u\Big(\gamma_\sigma(\omega)\Big) \\
&= \max_{\gamma \in F^S} \sum_{\omega \in \Omega} \pi(\omega) u\left(\sum_{s \in S} \sigma(s \mid \omega)[\gamma(s)](\omega)\right) \\
&= \max_{\gamma \in F^S} \sum_{\omega \in \Omega} \pi(\omega) \sum_{s \in S} \sigma(s \mid \omega) u\Big([\gamma(s)](\omega)\Big) \\
&= \sum_{s \in S} \max_{f \in F} \sum_{\omega \in \Omega} \pi(\omega) \sigma(s \mid \omega) u(f(\omega)) \\
&= W(\{F\}, \sigma)
\end{aligned}$$

where the second equality follows from the construction of F_σ, the third equality follows from the definition of an induced act γ_σ, the fourth equality follows from u being an affine function, and the fifth equality follows from the fact that maximizing over the set of all plans is equivalent to maximizing over the menu of acts contingent on each signal realization. Therefore, $\succsim$ must satisfy Reduction.

If. Suppose $\succsim$ satisfies Weak Order, Act Independence, SRNI and Reduction, and $\succsim^\sigma$ satisfies Axioms 1.2-1.8 for each $\sigma \in \mathcal{I}$, we want to show that $\succsim$ has a regret-varying IT representation.

As argued in the main text, if $\succsim$ satisfies Weak Order and Reduction, then $\succsim$ satisfies Stable Preference over Acts. Therefore, we can apply Lemma 1.4 to conclude that $\succsim$ has an aligned informational tradeoff representation with parameters $\left(\pi, u, \left(K^{\sigma}\right)_{\sigma\in\mathcal{I}}, \left(i^{\sigma}\right)_{\sigma\in\mathcal{I}}\right)$, we want to show that if $\succsim$ satisfies SRNI and Reduction, then i^{σ} and σ induce the same distribution over posteriors given π for any $\sigma \in \mathcal{I}$.

By Lemma 1.5, SRNI implies that i^{o} coincides with o, that is,

$$W(\mathbb{F}, o) = \max_{F\in\mathbb{F}} \max_{f\in F} \sum_{\omega\in\Omega} \pi(\omega)u(f(\omega)).$$

By Reduction, $W(\{F\}, \sigma) = W(\{F_{\sigma}\}, o)$ for any menu $F \in \mathcal{M}$ and information structure $\sigma \in \mathcal{I}$. Note that

$$W(\{F\}, \sigma) = \sum_{t\in S^{\sigma}} \max_{f\in F} \sum_{\omega\in\Omega} \pi(\omega)i^{\sigma}(t \mid \omega)u(f(\omega))$$

$$W(\{F_{\sigma}\}, o) = \sum_{s\in S} \max_{f\in F} \sum_{\omega\in\Omega} \pi(\omega)\sigma(s \mid \omega)u(f(\omega))$$

where the first equation follows from applying W to the pair $(\{F\}, \sigma)$ and S^{σ} is the set of signal realizations for the identified information structure i^{σ}, that is, $i^{\sigma} : \Omega \to \Delta(S^{\sigma})$. The second equation follows from the previous derivation in the proof for the necessity of Reduction (S is the set of signal realizations for σ and generally different from S^{σ}). Therefore, σ and i^{σ} must induce the same distribution over posteriors (otherwise, by a standard result, we can find a menu F such that $W(\{F\}, \sigma) \neq W(\{F_{\sigma}\}, o)$). This completes the proof. □

A.3.5 Proof of Theorem 1.3

Proof. Only if. Suppose $\succsim$ has an information tradeoff representation, then $\succsim$ has a regret-varying IT representation with $K^{\sigma} = K$ for all $\sigma \in \mathcal{I}$. Therefore, we can apply Lemma 1.6 to conclude that $\succsim$ satisfies Weak Order, Act Independence, SRNI and Reduction, and $\succsim^{\sigma}$ satisfies Axioms 1.2-1.8 for each $\sigma \in \mathcal{I}$.

We now check that $\succsim$ must satisfy Balance. Note that for any menu $F \in \mathcal{M}$ and any information structure $\sigma \in \mathcal{I}$,

$$W(D(F), \sigma) = (1 + K)W(\{F\}, o) - K \cdot W(\{F\}, \sigma). \quad \text{(A.23)}$$

Rearranging, we have

$$W(\{F\}, o) = \frac{1}{1+K} \cdot W(D(F), \sigma) + \frac{K}{1+K} \cdot W(\{F\}, \sigma). \quad \text{(A.24)}$$

Suppose $(\{F\}, \sigma) \succ (\{F\}, o)$, then $W(\{F\}, \sigma) > W(\{F\}, o)$, which further indicates that $W(\{F\}, \sigma) > W(D(F), \sigma)$, otherwise the above equation cannot hold.

Now further suppose that for some $\alpha \in (0, 1]$,

$$\Big(\alpha D(F) + (1 - \alpha)\{F\}, \sigma\Big) \sim (\{F\}, o).$$

Then

$$W(\{F\}, o) = \alpha \cdot \widetilde{W}(D(F), \sigma) + (1 - \alpha) \cdot W(\{F\}, \sigma). \quad \text{(A.25)}$$

For equations (A.24) and (A.25) to hold at the same time, it must be that $\alpha = \frac{1}{1+K}$. Then $\succsim$ must satisfy Balance by the fact that equation (A.24) is an identity that holds for all pairs $(F, \sigma) \in \mathcal{M} \times \mathcal{I}$.

If. Suppose $\succsim$ satisfies Weak Order, Act Independence, SRNI, Reduction and Balance, and $\succsim^\sigma$ satisfies Axioms 1.2-1.8 for each $\sigma \in \mathcal{I}$. We want to show that $\succsim$ has an IT representation.

First, we can apply Lemma 1.6 to conclude that $\succsim$ has a regret-varying IT representation with parameters $\big(\pi, u, (K^\sigma)_{\sigma \in \mathcal{I}}\big)$. We want to show that if $\succsim$ also satisfies Balance, then there exists $K \geq 0$ such that setting $K^\sigma = K$ for all $\sigma \in \mathcal{I}$ will deliver the same utility representation as W.

If some σ induces a degenerate distribution over posteriors, then we can set K^σ to be any non-negative scalar without affecting the value of W. So we only need to worry about information structures that induce non-degenerate distributions over posteriors.

Suppose by contradiction that $K^{\sigma} \neq K^{\sigma'}$ for some $\sigma, \sigma' \in \mathcal{I}$ such that σ and σ' each induces a non-degenerate distribution over posteriors. We want to show that Balance must be violated.

By assumption, we can find menus F and G such that $(\{F\}, \sigma) \succ (\{F\}, o)$ and $(\{G\}, \sigma') \succ (\{G\}, o)$. Let $\alpha = 1/(1 + K^{\sigma})$ and $\alpha' = 1/(1 + K^{\sigma'})$, then $\alpha \neq \alpha'$. Without loss of generality, suppose $\alpha > \alpha'$. Follow a similar computation in the proof for the necessity of Balance,

$$\Big(\alpha' D(F) + (1-\alpha')\{F\}, \sigma\Big) \succ \Big(\alpha D(F) + (1-\alpha)\{F\}, \sigma\Big) \sim (\{F\}, o)$$

$$\Big(\alpha' D(G) + (1-\alpha)\{G\}, \sigma'\Big) \sim (\{G\}, o)$$

which is a direct violation of Balance. This completes the proof of the sufficiency of the axioms for the existence of an IT representation.

The uniquenuss of an IT representation is an immediate corollary of the uniqueness of the SIT represenation for each $\succsim^{\sigma}$. □

A.3.6 Proof of Lemma 1.7

Proof. Combining equations (1.21) and (1.22), we see that

$$\widetilde{U}(F, \mathbb{F}, \mu_s^{\sigma}) = (1 + K_1 + K_2) \max_{f \in F} \sum_{\omega \in \Omega} \mu_s^{\sigma}(\omega) u(f(\omega))$$

$$- K_1 \sum_{\omega \in \Omega} \mu_s^{\sigma}(\omega) \max_{G \in \mathbb{F}} \max_{g \in F} u(g(\omega)) - K_2 \sum_{\omega \in \Omega} \mu_s^{\sigma}(\omega) \max_{h \in F} u(h(\omega)) \tag{A.26}$$

Plug this back to (1.20), we can write W_1 as

$$W_1(\mathbb{F},\sigma) = \max_{F\in\mathbb{F}}\Bigg[\sum_{s\in S}(1+K_0+K_1+K_2)\max_{f\in F}\sum_{\omega\in\Omega}\pi(\omega)\sigma(s\mid\omega)u(f(\omega)) - K_2\sum_{\omega\in\Omega}\pi(\omega)\sigma(s\mid\omega)\max_{h\in F}u(h(\omega))\Bigg] - K_0\sum_{s\in S}\max_{G\in\mathbb{F}}\max_{g\in G}\sum_{\omega\in\Omega}\pi(\omega)\sigma(s\mid\omega)u(g(\omega)) - K_1\sum_{\omega\in\Omega}\pi(\omega)\max_{G\in\mathbb{F}}\max_{g\in G}u(g(\omega)) \tag{A.27}$$

If $K_2 = 0$, this reduces to

$$W_1(\mathbb{F},\sigma) = (1+K_0+K_1)\max_{F\in\mathbb{F}}\sum_{s\in S}\max_{f\in F}\sum_{\omega\in\Omega}\pi(\omega)\sigma(s\mid\omega)u(f(\omega)) - K_0\sum_{s\in S}\max_{G\in\mathbb{F}}\max_{g\in G}\sum_{\omega\in\Omega}\pi(\omega)\sigma(s\mid\omega)u(g(\omega)) - K_1\sum_{\omega\in\Omega}\pi(\omega)\max_{G\in\mathbb{F}}\max_{g\in G}u(g(\omega)) \tag{A.28}$$

Note that the last term does not depend on σ. For convenience, let

$$U_1(\mathbb{F},\sigma) := \max_{F\in\mathbb{F}}\sum_{s\in S}\max_{f\in F}\sum_{\omega\in\Omega}\pi(\omega)\sigma(s\mid\omega)u(f(\omega))$$

$$U_2(\mathbb{F},\sigma) := \sum_{s\in S}\max_{G\in\mathbb{F}}\max_{g\in G}\sum_{\omega\in\Omega}\pi(\omega)\sigma(s\mid\omega)u(g(\omega))$$

Therefore,

$$\begin{aligned}
& W_1(\mathbb{F},\sigma) \geq W_1(\mathbb{F},\sigma') \\
\iff & (1+K_0+K_1)\cdot U_1(\mathbb{F},\sigma) - K_0\cdot U_2(\mathbb{F},\sigma) \geq (1+K_0+K_1)\cdot U_1(\mathbb{F},\sigma') - K_0\cdot U_2(\mathbb{F},\sigma') \\
\iff & \frac{1+K_0+K_1}{1+K_1}\cdot U_1(\mathbb{F},\sigma) - \frac{K_0}{1+K_1}\cdot U_2(\mathbb{F},\sigma) \geq \frac{1+K_0+K_1}{1+K_1}\cdot U_1(\mathbb{F},\sigma') - \frac{K_0}{1+K_1}\cdot U_2(\mathbb{F},\sigma') \\
\iff & (1+K)\cdot U_1(\mathbb{F},\sigma) - K\cdot U_2(\mathbb{F},\sigma) \geq (1+K)\cdot U_1(\mathbb{F},\sigma') - K\cdot U_2(\mathbb{F},\sigma') \\
\iff & W(\mathbb{F},\sigma) \geq W(\mathbb{F},\sigma')
\end{aligned}$$

where the third equivalence follows from the assumption that $K = K_0/(1+K_1)$. □

A.4 Translation from Lotteries to Acts

A.4.1 Proof of Lemma A.20

Proof. We can check that the operation $u(\cdot)$ that maps a primitive (acts, menus or directions) into its corresponding primitive as in utilities (utility acts, utility menus and utility directions) behave nicely with set operations and respects linearity. This guarantees that each Axiom from 1.1-1.6 and 1.8 implies their counterpart in Axiom A.2.1-A.2.6 and A.8.

Note that the statement of Axiom A.2.7 is slightly different from the statement of Axiom 1.7. We now check that Axioms 1.1-8 implies Axiom A.2.7.

First notice that by Axiom 1.7, there exists lotteries $\ell, \ell' \in \Delta(X)$ such that $\ell \succ \ell'$. Since these are constant acts, it must be that ℓ dominates ℓ'. Then by Axiom 1.8, $\{\{\ell\}, \{\ell'\}\} \sim \{\{\ell\}\}$. Then by transitivity, $\{\{\ell\}, \{\ell'\}\} \succ \{\{\ell'\}\}$. Let $\mathbb{F} = \{\{\ell\}, \{\ell'\}\}$ and $\mathbb{G} = \{\{\ell'\}\}$. Then $\mathbb{F} \supseteq \mathbb{G}$ and $\mathbb{F} \succ \mathbb{G}$.

Then we can apply the definition of $\succsim_u$ to conclude that $u(\mathbb{F}), u(\mathbb{G}) \in \mathcal{D}_u$, and $u(\mathbb{F}) \supseteq u(\mathbb{G})$ with $u(\mathbb{F}) \succ_u u(\mathbb{G})$, Axiom A.2.7 is satisfied. □

A.4.2 Proof of Lemma A.21

Proof that Axiom A.2.1 implies Axiom A.2.9.

Fix any $\mathbb{F}', \mathbb{G}' \in \mathcal{D}'$. Then by the completeness of $\succsim_u$, either $r(\mathbb{F}') \succsim_u r(\mathbb{G}')$ or $r(\mathbb{G}') \succsim_u r(\mathbb{F}')$ or both, which further indicates that either $\mathbb{F}' \succsim_* \mathbb{G}'$ or $\mathbb{G}' \succsim_* \mathbb{F}'$ or both.

For transitivity, suppose $\mathbb{F}' \succsim_* \mathbb{G}'$ and $\mathbb{G}' \succsim_* \mathbb{H}'$, then by definition, $r(\mathbb{F}') \succsim_u r(\mathbb{G}')$ and $r(\mathbb{G}') \succsim_u r(\mathbb{H}')$, which implies $r(\mathbb{F}') \succsim_u r(\mathbb{H}')$ (by the transitivity of $\succsim_u$), which further implies that $\mathbb{F}' \succsim_* \mathbb{H}'$. □

Proof that Axiom A.2.2 implies Axiom A.2.10.

Fix any $\mathbb{F}' \in \mathcal{D}'$, consider the set $\{\mathbb{G}' \in \mathcal{D}' \mid \mathbb{F}' \succsim_* \mathbb{G}'\}$. By Axiom A.2.2, the set

$\{\mathbb{G}_u \in \mathcal{D}_u \mid r(\mathbb{F}') \succsim_u \mathbb{G}\}$ is closed. Since r is a homeomorphism, it suffices to show that

$$\{\mathbb{G}' \in \mathcal{D}' \mid \mathbb{F}' \succsim_* \mathbb{G}'\} = r^{-1}\Big(\{\mathbb{G}_u \in \mathcal{D}_u \mid r(\mathbb{F}') \succsim_u \mathbb{G}_u\}\Big).$$

If $\mathbb{F}' \succsim_* \mathbb{G}'$, then $r(\mathbb{F}') \succsim_u r(\mathbb{G}')$, thus LHS $\subseteq$ RHS. If $r(\mathbb{F}') \succsim_u \mathbb{G}_u$, then $\mathbb{F}' \succsim_* r^{-1}(\mathbb{G}_u)$, thus LHS $\supseteq$ RHS. The arguments for $\{\mathbb{G}' \in \mathcal{D}' \mid \mathbb{G}' \succsim_* \mathbb{F}'\}$ are similar. □

Proof that Axiom A.2.3 implies Axiom A.2.11.

First note that r is linear, that is, for any $\mathbb{F}', \mathbb{G}' \in \mathcal{D}'$ and $\alpha \in [0,1]$,

$$r(\alpha\mathbb{F}' + (1-\alpha)\mathbb{G}') = \alpha r(\mathbb{F}') + (1-\alpha) r(\mathbb{G}').$$

Then notice that r^{-1} is also linear. For any $\mathbb{F}, \mathbb{G} \in \mathcal{D}$ and $\alpha \in [0,1]$,

$$r^{-1}(\alpha\mathbb{F} + (1-\alpha)\mathbb{G}) = \alpha r^{-1}(\mathbb{F}) + (1-\alpha) r^{-1}(\mathbb{G}).$$

Therefore, for any $\mathbb{F}', \mathbb{G}', \mathbb{H}' \in \mathcal{D}'$ and any $\alpha \in (0,1]$,

$$\begin{aligned} \mathbb{F}' \succsim_* \mathbb{G}' &\iff r(\mathbb{F}') \succsim_u r(\mathbb{G}') \\ &\iff \alpha r(\mathbb{F}') + (1-\alpha) r(\mathbb{H}') \succsim_u \alpha r(\mathbb{G}') + (1-\alpha) r(\mathbb{H}') \\ &\iff r(\alpha\mathbb{F}' + (1-\alpha)\mathbb{H}') \succsim_u r(\alpha\mathbb{G}' + (1-\alpha)\mathbb{H}') \\ &\iff \alpha\mathbb{F}' + (1-\alpha)\mathbb{H}' \succsim_* \alpha\mathbb{G}' + (1-\alpha)\mathbb{H}' \end{aligned}$$

This is essentially the same proof as the proof of Claim 3 in DLST (2014). □

Proof that Axiom A.2.4 implies Axiom A.2.12.

Given $\succsim_u$, let N be a natural number satisfying Axiom A.2.4.

Fix any $\mathbb{F}' \in \mathcal{D}'$, by Axiom A.2.4, there exists $\mathbb{G}_u \in \mathcal{D}_u$ such that $\mathbb{G}_u$ is critical for $r(\mathbb{F}')$ with $|\mathbb{G}_u| < N$. Let $\mathbb{G}' = r^{-1}(\mathbb{G}_u)$. Then $|\mathbb{G}'| = |\mathbb{G}_u| < N$, and for any $\mathbb{H}'$ satisfying $\mathbb{G}' \subseteq \mathbb{H}' \subseteq \mathbb{F}'$, we have

$$r(\mathbb{G}') \subseteq r(\mathbb{H}') \subseteq r(\mathbb{F}') \Longrightarrow r(\mathbb{G}') \sim_u r(\mathbb{H}') \sim_u r(\mathbb{F}') \Longrightarrow G' \sim_* \mathbb{H}' \sim_* \mathbb{F}'.$$

Hence $\mathbb{G}'$ is critical for $\mathbb{F}'$ with $|\mathbb{G}'| < N$. Similar arguments for the second half. □

Proof that Axiom A.2.5 implies Axiom A.2.13.

First notice that for any $F' \in \mathcal{M}'$, $r(\{F'\}) = \{r(F')\}$ by construction. Therefore,

$$\begin{aligned}\{F'\} \succsim_* \{G'\}, F' \in \mathbb{F}' &\Longrightarrow \{r(F')\} \succsim_u \{r(G')\}, r(F') \in r(\mathbb{F}') \\ &\Longrightarrow r(\mathbb{F}') \succsim_u r(\mathbb{F}') \cup \{r(G')\} \\ &\Longrightarrow r(\mathbb{F}') \succsim_u r(\mathbb{F}' \cup \{G'\}) \Longrightarrow \mathbb{F}' \succsim_* \mathbb{F}' \cup \{G'\}\end{aligned}$$

This completes the proof. □

Proof that Axiom A.2.6 implies Axiom A.2.14.

First note that

$$r(\mathbb{F}' \cup \{F' \cup G'\}) = r(\mathbb{F}') \cup \{r(F') \cup r(G')\}$$

and that

$$r(\mathbb{F}' \cup \{F', G'\}) = r(\mathbb{F}') \cup \{r(F'), r(G')\}.$$

Since $r(\mathbb{F}') \in \mathcal{D}_u$, and $r(F'), r(G') \in \mathcal{M}_u$, then we can use Axiom A.2.6 to conclude that

$$r(\mathbb{F}') \cup \{r(F') \cup r(G')\} \succsim_u r(\mathbb{F}') \cup \{r(F'), r(G')\},$$

which further indicates that $\mathbb{F}' \cup \{F' \cup G'\} \succsim_* \mathbb{F}' \cup \{F', G'\}$. □

Proof that Axiom A.2.7 implies Axiom A.2.15.

By Axiom A.2.7, there exists $\mathbb{F}_u, \mathbb{G}_u \in \mathcal{D}_u$ such that $\mathbb{F}_u \supseteq \mathbb{G}_u$ and $\mathbb{F}_u \succ_u \mathbb{G}_u$. Therefore, $r^{-1}(\mathbb{F}_u), r^{-1}(\mathbb{G}_u) \in \mathcal{D}'$ satisfy $r^{-1}(\mathbb{F}_u) \supseteq r^{-1}(\mathbb{G}_u)$ and $r^{-1}(\mathbb{F}_u) \succ_* r^{-1}(\mathbb{G}_u)$.

□

Proof that Axiom A.2.8 implies Axiom A.2.16.

Suppose $G' \subseteq F'$, then $r(G') \subseteq r(F')$, and by the definition of domination, $r(F')$ dominates $r(G')$. Suppose $F' \in \mathbb{F}'$, then $r(F') \in r(\mathbb{F}')$, and the second part of Axiom A.2.8 implies that $r(\mathbb{F}') \sim_u r(\mathbb{F}') \cup \{r(G')\}$, which further indicates that $\mathbb{F}' \sim_* \mathbb{F}' \cup \{G'\}$. □

A.4.3 Proof of Lemma A.22

Before we move on to study the properties of $\succsim_{**}$, we have the following simplifying result to help up use the definition of $\succsim_{**}$ more easily.

Lemma A.26. *Suppose $\succsim_u$ satisfies Axioms A.2.1-A.2.3 (Weak Order, Continuity and Independence), then $\succsim_{**}$ constructed as in Step 2 will satisfy*

$$\mathbb{F}'' \succsim_{**} \mathbb{G}'' \iff \frac{1}{n^2}\mathbb{F}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\} \succsim_{*} \frac{1}{n^2}\mathbb{G}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\}.$$

Proof of Lemma A.26. Since $\succsim_u$ satisfies Axioms A.2.1-A.2.3, then by Lemma A.21, $\succsim_*$ satisfies Axioms A.2.9-A.2.11. Recall that

$$F^{n+1} = \left\{\left(\frac{n}{n+1}, \dots, \frac{n}{n+1}\right)\right\}.$$

That is, F^{n+1} is the singleton menu containing only a "uniform lottery."

"$\Longleftarrow$:" Suppose $\frac{1}{n^2}\mathbb{F}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\} \succsim_* \frac{1}{n^2}\mathbb{G}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\}$. We know that $\{F^{n+1}\} \in \mathcal{D}'$ and $\succsim_*$ satisfies A.2.11. Therefore, for any $\varepsilon \in [0, 1/n^2)$, $\alpha := n^2\varepsilon \in [0, 1)$ and

$$\begin{aligned}
\varepsilon\mathbb{F}'' + (1-\varepsilon)\{F^{n+1}\} &= \alpha\left(\frac{1}{n^2}\mathbb{F}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\}\right) + (1-\alpha)\{F^{n+1}\} \\
&\succsim_* \alpha\left(\frac{1}{n^2}\mathbb{G}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\}\right) + (1-\alpha)\{F^{n+1}\} \\
&= \varepsilon\mathbb{G}'' + (1-\varepsilon)\{F^{n+1}\}
\end{aligned}$$

"$\Longrightarrow$:" Suppose $\mathbb{F}'' \succsim_{**} \mathbb{G}''$, then we can find a sequence of ε_n converging from below to $1/n^2$, then the result follows from the continuity of $\succsim_*$. □

An immediate but important implication of Lemma A.26 is that the constructed binary relation $\succsim_{**}$ is uniquely determined by $\succsim_*$.

Moving forward, we will do the operation "mix $\mathbb{F}''$ with $\{F^{n+1}\}$ on weight $1/n^2$" a lot in the following proofs. For an easier exposition, we give this operation a name by defining $t : \mathcal{D}'' \to \mathcal{D}'$ given by

$$t(\mathbb{F}'') := \frac{1}{n^2}\mathbb{F}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\} \tag{A.29}$$

Proof that Axioms A.2.9-A.2.11 imply Axiom A.2.17.

Completeness: Fix any $\mathbb{F}'', \mathbb{G}'' \in \mathcal{D}''$, $t(\mathbb{F}''), t(\mathbb{G}'') \in \mathcal{D}'$. By the completeness of $\succsim_*$, $t(\mathbb{F}') \succsim_* t(\mathbb{G}')$ or $t(\mathbb{G}') \succsim_* t(\mathbb{F}')$ or both. By Lemma A.26, Axioms A.2.9-A.2.11 guarantee that $\mathbb{F}'' \succsim_{**} \mathbb{G}''$ if and only if $t(\mathbb{F}'') \succsim_* t(\mathbb{G}'')$. Thus, $\succsim_{**}$ is complete.

Transitivity:

$$\begin{aligned} \mathbb{F}'' \succsim_{**} \mathbb{G}'',\ \mathbb{G}'' \succsim_{**} \mathbb{H}'' &\Longrightarrow t(\mathbb{F}'') \succsim_* t(\mathbb{G}''),\ t(\mathbb{G}'') \succsim_* t(\mathbb{H}'') \\ &\Longrightarrow t(\mathbb{F}'') \succsim_* t(\mathbb{H}'') \\ &\Longrightarrow \mathbb{F}'' \succsim_{**} \mathbb{H}'' \end{aligned}$$

where the second implication follows from $\succsim_*$ being transitive. □

Proof that Axioms Axioms A.2.9-A.2.11 imply Axiom A.2.18.

We consider the lower contour set. Same arguments work for the upper contour set.

Fix any sequence $\{\mathbb{G}''_m\}_{m=1}^{\infty} \subset \{\mathbb{G}'' \mid \mathbb{F}'' \succsim_{**} \mathbb{G}''\}$, that is, $\mathbb{F}'' \succsim_{**} \mathbb{G}''_m$ for all $m \in \mathbb{N}$, which further indicates that $t(\mathbb{F}'') \succsim_* t(\mathbb{G}''_m)$ for all $m \in \mathbb{N}$. If $\mathbb{G}''_m \to \mathbb{G}''_*$ as $m \to \infty$ for some $\mathbb{G}''_* \in \mathcal{D}''$, then by construction, $t(\mathbb{G}''_m) \to t(\mathbb{G}''_*)$. Since $\succsim_*$ satisfies Continuity (Axiom A.2.10), this implies that $t(\mathbb{F}'') \succsim_* t(\mathbb{G}''_*)$. Thus, $\mathbb{F}'' \succsim_{**} \mathbb{G}''_*$. □

Proof that Axioms A.2.9-A.2.11 imply Axiom A.2.19.

First note that the mapping $t : \mathcal{D}'' \to \mathcal{D}'$ is linear, that is, for any $\alpha \in [0,1]$,

$$\begin{aligned} t\Big(\alpha\mathbb{F}'' + (1-\alpha)\mathbb{G}''\Big) &= \frac{1}{n^2}\Big(\alpha\mathbb{F}'' + (1-\alpha)\mathbb{G}''\Big) + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\} \\ &= \frac{1}{n^2}\Big(\alpha\mathbb{F}'' + (1-\alpha)\mathbb{G}''\Big) + \left(1 - \frac{1}{n^2}\right)\Big(\alpha\{F^{n+1}\} + (1-\alpha)\{F^{n+1}\}\Big) \\ &= \alpha t(\mathbb{F}'') + (1-\alpha)t(\mathbb{G}'') \end{aligned}$$

Note that the second equality only goes through because F^{n+1} is a singleton menu. The decomposition is not generally doable for non-singleton menus.

Using the linearity of t, we prove the following stronger version of Independence. For any $\mathbb{F}'', \mathbb{G}'', \mathbb{H}''$ and any $\alpha \in (0,1]$,

$$\begin{aligned} \mathbb{F}'' \succsim_{**} \mathbb{G}'' &\iff t(\mathbb{F}'') \succsim_{*} t(\mathbb{G}'') \\ &\iff \alpha t(\mathbb{F}'') + (1-\alpha)t(\mathbb{H}'') \succsim_{*} \alpha t(\mathbb{G}'') + (1-\alpha)t(\mathbb{H}'') \\ &\iff t\Big(\alpha\mathbb{F}'' + (1-\alpha)\mathbb{H}''\Big) \succsim_{*} t\Big(\alpha\mathbb{G}'' + (1-\alpha)\mathbb{H}''\Big) \\ &\iff \alpha\mathbb{F}'' + (1-\alpha)\mathbb{H}'' \succsim_{**} \alpha\mathbb{G}'' + (1-\alpha)\mathbb{H}'' \end{aligned}$$

where the first and the last equivalence follow from Lemma A.26, the second equivalence follows from $\succsim_*$ satisfying Axiom A.2.11 and the third equivalence follows from the linearity of t we showed above. □

Proof that Axioms A.2.9-A.2.12 imply Axiom A.2.20.

First note that the mapping t respects set inclusion, that is, if $\mathbb{G}'' \subseteq \mathbb{F}''$, then

$$t(\mathbb{G}'') = \frac{1}{n^2}\mathbb{G}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\} \subseteq \frac{1}{n^2}\mathbb{F}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\} = t(\mathbb{F}'').$$

Moreover, $|t(\mathbb{F}'')| = |\mathbb{F}''|$ for any finite $\mathbb{F}''$.

Now Axiom A.2.4 implies Axiom A.2.12, which guarantees the existence of a natural number N satisfying the conditions for $\succsim_*$ over $\mathcal{D}'$ corresponding to the two conditions above. We argue the same N works for $\succsim_{**}$.

Fix any $\mathbb{F}'' \in \mathcal{D}''$, $t(\mathbb{F}'') \in \mathcal{D}'$, and thus by Axiom A.2.12, there exists $\mathbb{G}'$ such that $\mathbb{G}'$ is critical for $t(\mathbb{F}'')$ and $|\mathbb{G}'| < N$. Since $\mathbb{G}' \subseteq t(\mathbb{F}'')$, there exists $\mathbb{G}'' \subseteq \mathbb{F}''$ such that $\mathbb{G}' = t(\mathbb{G}'')$. So $|\mathbb{G}''| = |\mathbb{G}'| < N$, and we argue that $\mathbb{G}''$ is critical for $\mathbb{F}''$. Fix any $\mathbb{H}''$ such that $\mathbb{G}'' \subseteq \mathbb{H}'' \subseteq \mathbb{F}''$, then $\mathbb{G}' = t(\mathbb{G}'') \subseteq t(\mathbb{H}'') \subseteq t(\mathbb{F}'')$. This implies that $G' \sim_* t(\mathbb{H}'') \sim_* t(\mathbb{F}'')$ since $\mathbb{G}'$ is critical for $t(\mathbb{F}'')$, which further implies that $\mathbb{H}'' \sim_{**} \mathbb{F}''$.

Similar arguments work for the second condition. □

Proof that Axioms A.2.9-A.2.11 and A.2.13 imply Axiom A.2.21.

First note that the mapping $t : \mathcal{D}'' \to \mathcal{D}'$ respects set unions: For any $\mathbb{F}'', \mathbb{G}'' \in \mathcal{D}''$,

$$\begin{aligned} t\Big(\mathbb{F}'' \cup \mathbb{G}''\Big) &= \frac{1}{n^2}\Big(\mathbb{F}'' \cup \mathbb{G}''\Big) + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\} \\ &= \left(\frac{1}{n^2}\mathbb{F}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\}\right) \cup \left(\frac{1}{n^2}\mathbb{G}'' + \left(1 - \frac{1}{n^2}\right)\{F^{n+1}\}\right) \\ &= t(\mathbb{F}'') \cup t(\mathbb{G}'') \end{aligned}$$

Now suppose $\mathbb{F}''$ and $\mathbb{G}''$ satisfy that for any $G'' \in \mathbb{G}''$, there exists $F'' \in \mathbb{F}''$ such that $\{F''\} \succsim_{**} \{G''\}$, we want to show that $\mathbb{F}'' \succsim_{**} \mathbb{F}'' \cup \mathbb{G}''$.

For any $G' \in t(\mathbb{G}'')$, there exists $G'' \in \mathbb{G}''$ such that $\{G'\} = t(\{G''\})$. By assumption, there exists $F'' \in \mathbb{F}''$ such that $\{F''\} \succsim_{**} \{G''\}$. Then $t(\{F''\}) \succsim_* t(\{G''\}) = \{G'\}$ and $t(F'') \in t(\mathbb{F}'')$. (This is an abuse of notation, by $t(F'')$ we mean the only menu contained in $t(\{F''\})$.) Then the assumption in Axiom A.2.13 is satisfied and we can conclude that $t(\mathbb{F}'') \succsim_* t(\mathbb{F}'') \cup t(\mathbb{G}'') = t(\mathbb{F}'' \cup \mathbb{G}'')$, which further indicates that $\mathbb{F}'' \succsim_{**} \mathbb{F}'' \cup \mathbb{G}''$. □

Proof that Axioms A.2.9-A.2.11 and A.2.14 imply Axiom A.2.22.

We first formalize the notation (abuse) appeared above. Let $F'' \in \mathcal{M}''$ be a menu, then

$$t(F'') := \frac{1}{n^2}F'' + \left(1 - \frac{1}{n^2}\right)F^{n+1} \in \mathcal{M}'.$$

And $t(F'' \cup G'') = t(F'') \cup t(G'')$ for all $F'', G'' \in \mathcal{M}''$. Now for any $\mathbb{F}'' \in \mathcal{D}''$ and any $F'', G'' \in \mathcal{M}''$, $t(\mathbb{F}'') \in \mathcal{D}'$, $t(F''), t(G'') \in \mathcal{M}'$. Thus, by Axiom A.2.14,

$$t(\mathbb{F}'') \cup \{t(F'') \cup t(G'')\} \succsim_* t(\mathbb{F}'') \cup \{t(F''), t(G'')\}.$$

But $t\Big(\mathbb{F}'' \cup \{F'' \cup G''\}\Big) = t(\mathbb{F}'') \cup \{t(F'' \cup G'')\} = t(\mathbb{F}'') \cup \{t(F'') \cup t(G'')\}$, and similar for the RHS. Therefore, $\mathbb{F}'' \cup \{F'' \cup G''\} \succsim_{**} \mathbb{F}'' \cup \{F'', G''\}$. $\square$

Proof that Axiom A.2.9-A.2.11 and A.2.15 imply Axiom A.2.23.

If $F'' \supseteq G''$, then $t(F'') \supseteq t(G'')$, and by Axiom 7*, $t(\mathbb{F}'') \cup \{t(F''), t(G'')\} \succsim_* t(\mathbb{F}'') \cup \{t(F'')\}$. $\square$

Appendix B

Appendix to Chapter 2

B.1 General Auxiliary State Space

In this section, we consider auxiliary state spaces beyond finite sets and we will show that our results in the main text remain valid in this more general environment.

Let the auxiliary state space Θ be an arbitrary nonempty set. Ambiguous experiments are defined the same way as in Definition 2.4, that is, an *ambiguous experiment* is a mapping $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$, where the payoff-relevant space Ω and the set of signal realizations S are still assumed to be finite.

Consider an individual with a finite set of actions A, a state-dependent utility function $u : \Omega \times A \to \mathbb{R}$ and a prior belief $\pi \in \Delta(\Omega)$. The *expected utility conditional on* state θ, $U\big(\sigma, \mathbf{p}(\cdot, \theta)\big)$, is defined the same way as in equation (2.2):

$$U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) = \sum_{s \in S} \sum_{\omega \in \Omega} \sum_{a \in A} \pi(\omega)\mathbf{p}(s \mid \omega, \theta)\sigma(a \mid s)u(\omega, a).$$

Note that U is bounded as a function of θ for any fixed $(\mathbf{p}, A, u, \pi, \sigma)$. We say $V : \mathbb{R}^\Theta \to \mathbb{R}$ is a *monotone aggregator* if $V(f) \geq V(g)$ whenever two functions $f, g : \Theta \to \mathbb{R}$ satisfy $f(\theta) \geq g(\theta)$ for all $\theta \in \Theta$. Comparing to Definition 2.5, we drop the continuity requirement on the aggregator. Let $\mathcal{V}_{Mono}$ denote the set of all monotone aggregators.

Let 2^Θ denote the set of all subsets of Θ. Let Δ denote the set of all finitely additive probability measures over 2^Θ. Fixing any $\mu \in \Delta$, let V_μ denote the aggregator that corresponds to the (subjective) expected utility index, that is,

$$V_\mu \left(\left\{ U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) \right\}_{\theta \in \Theta} \right) := \int_\Theta U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) \ d\mu(\theta). \tag{B.1}$$

The integral is well-defined as $U\big(\sigma, \mathbf{p}(\cdot, \theta)\big)$ is a bounded measurable function and μ is a finitely additive measure (Aliprantis and Border, 2006, Theorem 11.8). Let $\mathcal{V}_{SEU}$ denote the class of all aggregators that corresponds to some (subjective) expected utility index, that is, $\mathcal{V}_{SEU} := \{V_\mu : \mathbb{R}^\Theta \to \mathbb{R} \mid \mu \in \Delta\}$. When Θ is uncountable, $\mathcal{V}_{SEU}$ includes the set of all non-atomic (finitely additive) probability measures, and thus includes the standard subjective expected utility model à la Savage (1954).

As we no longer require any continuity of V, we slightly modify the definition of the ex-ante utility for an ambiguous experiment. Fixing an ambiguous experiment $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ and (A, u, π, V), the individual's *ex-ante utility* is

$$\sup_{\sigma \in A_S} V \left(\left\{ U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) \right\}_{\theta \in \Theta} \right).$$

Comparing to the definition of the ex-ante utility in the main text, we have the supremum operator instead of the maximum operator. This change has no impact on the interpretation of our comparison result: If $\mathbf{p}$ gives higher ex-ante utility than $\mathbf{p}'$, then for any action plan σ' that can be made when the DM faces $\mathbf{p}'$, he can find another action plan σ that guarantees

$$V \left(\left\{ U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) \right\}_{\theta \in \Theta} \right) \geq V \left(\left\{ U\big(\sigma', \mathbf{p}'(\cdot, \theta)\big) \right\}_{\theta \in \Theta} \right).$$

Moreover, without any continuity requirement, we can include more aggregators that corresponds to more ambiguity preferences.

For any $\mu \in \Delta$, the *expected experiment* with respect to μ, $\mathbf{p}_\mu : \Omega \to \Delta(S)$, is a Blackwell experiment defined similarly as in Definition 2.7 by

$$\mathbf{p}_\mu(s \mid \omega) := \int_\Theta \mathbf{p}(s \mid \omega, \theta)\ d\mu(\theta),\ \forall (s, \omega) \in S \times \Omega.$$

Let Δ_0 denote the set of all simple probability measures (i.e., probability measures with finite supports) over 2^Θ. For any $\mu \in \Delta_0$, the expected experiment $\mathbf{p}_\mu$ is simply

$$\mathbf{p}_\mu(s \mid \omega) = \sum_{\theta \in \Theta} \mathbf{p}(s \mid \omega, \theta)\mu(\theta),\ \forall (s, \omega) \in S \times \Omega.$$

We say that an ambiguous experiment $\mathbf{p} : \Omega \times \Theta \to \Delta(S)$ *prior-by-prior dominates* another $\mathbf{p}' : \Omega \times \Theta \to \Delta(S')$ if $\mathbf{p}_\mu$ is Blackwell more informative than $\mathbf{p}'_\mu$ for any $\mu \in \Delta_0$. Note that when Θ is finite, this coincides with our definition of prior-by-prior dominance in the main text (Definition 2.8). Finally, let $\mathcal{V}_0 := \{V_\mu \mid \mu \in \Delta_0\}$ where $V_\mu : \mathbb{R}^\Theta \to \mathbb{R}$ is defined as in equation (B.1).

Theorem B.1. *Let $p : \Omega \times \Theta \to \Delta(S)$ and $p' : \Omega \times \Theta \to \Delta(S')$ be two ambiguous experiments and let $\mathcal{V}$ be a class of aggregators such that $\mathcal{V}_0 \subset \mathcal{V} \subset \mathcal{V}_{Mono}$, then the following statements are equivalent:*

1. $\mathbf{p}$ *prior-by-prior dominates* $\mathbf{p}'$.

2. *For any A, u, π and any $\sigma' \in A_{S'}$, there exists $\sigma \in A_S$ such that*

$$U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) \geq U\big(\sigma', \mathbf{p}'(\cdot, \theta)\big) \text{ for all } \theta \in \Theta.$$

3. $\mathbf{p}$ *is preferred to* $\mathbf{p}'$ *in every decision problem by every decision maker whose ambiguity preferences can be represented by some $V \in \mathcal{V}$. That is,* $\mathbf{p}$ *gives weakly higher ex-ante utility than* $\mathbf{p}'$ *for every A, u, π and every $V \in \mathcal{V}$.*

Proof. See Appendix B.2.7. □

When Θ is finite, $\mathcal{V}_0$ reduces to $\mathcal{V}_{EU}$ as defined in equation (2.3.1) in the main text and Theorem B.1 coincides with Theorem 2.2. In addition to this characterization result, we also want to establish the equivalence of comparisons of ambiguous experiments and comparisons of sets of Blackwell experiments in this more general environment. Consider the Wald aggregator in this more general setting. Formally, let $V_W : \mathbb{R}^\Theta \to \mathbb{R}$ be re-defined by

$$V_W(f) := \inf_{\theta\in\Theta} f(\theta), \tag{B.2}$$

and we apply this aggregator to $f(\theta) = U\big(\sigma, \mathbf{p}(\cdot,\theta)\big)$. Let sets of experiments be defined in the same way as in Section 4.1, that is, a pair (S, P) where S is a finite set of signal realizations and P is a closed set (under the Euclidean topology) of Blackwell experiments, and P could be uncountable. Then we have the following proposition corresponding to Proposition 2.7 in the main text.

Proposition B.2. *For any two sets of Blackwell experiments* (S, P) *and* (S', P')*, we can construct an auxiliary state space* Θ*, and two ambiguous experiments* $\mathbf{p} : \Omega\times\Theta \to \Delta(S)$ *and* $\mathbf{p}' : \Omega\times\Theta \to \Delta(S')$ *as in* (2.14)*. Then for all* A*,* u*,* π*,* σ *and* σ'*:*

$$V_W\left(\left\{U\big(\sigma,\mathbf{p}(\cdot,\theta)\big)\right\}_{\theta\in\Theta}\right) = \min_{p\in P} U(\sigma,p),\ V_W\left(\left\{U\big(\sigma',\mathbf{p}'(\cdot,\theta)\big)\right\}_{\theta\in\Theta}\right) = \min_{p'\in P'} U(\sigma',p')$$

where V_W *is the Wald aggregator as defined in equation* (B.2)*, and* $U\big(\sigma,\mathbf{p}(\cdot,\theta)\big)$ *is the conditional expected utility defined in equation* (2.2)*.*

Proof. The proof is essentially the same with that of Proposition 2.7. The infimum on the left hand side can be replaced by a minimum since the $\Theta = P \times P'$ is compact and our construction of $\mathbf{p}$ and $\mathbf{p}'$ guarantees the continuity of $U\big(\sigma,\mathbf{p}(\cdot,\theta)\big)$ in θ (this is because $\theta \mapsto \mathbf{p}(\cdot,\theta)$ is just $(p,p') \mapsto p$, which can be viewed as a projection map and thus continuous as we endow $P \times P'$ with the product topology). $\square$

B.2 Proofs

B.2.1 Proof of Proposition 2.3

Proof. Suppose $\mathbf{p}$ globally Blackwell dominates $\mathbf{p}'$, then there exists a garbling $\gamma : S \to \Delta(S')$ such that $\mathbf{p}'(\cdot, \theta) = \gamma \circ \mathbf{p}(\cdot, \theta)$ for all $\theta \in \Theta$. Then we must have $\mathbf{p}'_\mu = \gamma \circ \mathbf{p}_\mu$ for any $\mu \in \Delta(\Theta)$, since for any $(s', \omega) \in S' \times \Omega$,

$$\begin{aligned}\mathbf{p}'_\mu(s' \mid \omega) &= \sum_{\theta \in \Theta} \mathbf{p}'(s' \mid \omega, \theta)\mu(\theta) \\ &= \sum_{\theta \in \Theta} \sum_{s \in S} \gamma(s' \mid s)\mathbf{p}(s \mid \omega, \theta)\mu(\theta) \\ &= \sum_{s \in S} \gamma(s' \mid s) \sum_{\theta \in \Theta} \mathbf{p}(s \mid \omega, \theta)\mu(\theta) = \sum_{s \in S} \gamma(s' \mid s)\mathbf{p}_\mu(s \mid \omega)\end{aligned}$$

To prove that the converse is not true, consider the following example:

$\Omega = \{\omega_1, \omega_2\}$, $\Theta = \{\theta_1, \theta_2\}$, $S = \{s_1, s_2\}$, $S' = \{s'_1, s'_2\}$, and

$\mathbf{p}(\cdot, \theta_1) =$

	s_1	s_2
ω_1	1	0
ω_2	0	1

$\mathbf{p}(\cdot, \theta_2) =$

	s_1	s_2
ω_1	0.8	0.2
ω_2	0.2	0.8

$\mathbf{p}'(\cdot, \theta_1) =$

	s'_1	s'_2
ω_1	0.8	0.2
ω_2	0.2	0.8

$\mathbf{p}'(\cdot, \theta_2) =$

	s'_1	s'_2
ω_1	0.8	0.2
ω_2	0.2	0.8

Then any $\mu \in \Delta(\Theta)$ can be represented by one parameter $k \in [0, 1]$, and the expected experiments are given by $\mathbf{p}_k = k\mathbf{p}(\cdot, \theta_1) + (1 - k)\mathbf{p}(\cdot, \theta_2)$. Thus,

$\mathbf{p}_k =$

	s_1	s_2
ω_1	$0.8 + 0.2k$	$0.2 - 0.2k$
ω_2	$0.2 - 0.2k$	$0.8 + 0.2k$

and $\mathbf{p}'_k =$

	s'_1	s'_2
ω_1	0.8	0.2
ω_2	0.2	0.8

Hence $\mathbf{p}'_k = \gamma_k \circ \mathbf{p}_k$ for any $k \in [0, 1]$ and the corresponding garbling γ_k for each k is

$\gamma_k =$

	s'_1	s'_2
s_1	$\frac{3+k}{3+2k}$	$\frac{k}{3+2k}$
s_2	$\frac{k}{3+2k}$	$\frac{3+k}{3+2k}$

$\mathbf{p}$ does not globally Blackwell dominates $\mathbf{p}'$ since γ_k varies with the belief k. □

B.2.2 Proof of Lemma 2.4

Proof of Lemma 2.4. $\mathbf{p}$ prior-by-prior dominates $\mathbf{q}$ if $\mathbf{p}_\eta$ is Blackwell more informative than $\mathbf{q}_\eta$ for every $\eta \in \Delta(\Theta_1 \times \Theta_2)$. Fix some $\eta \in \Delta(\Theta_1 \times \Theta_2)$ and let η_1 and η_2 be the marginal distributions over Θ_1 and Θ_2 induced by η, respectively.

$$\mathbf{p}_\eta = \sum_{(\theta_1,\theta_2)\in\Theta_1\times\Theta_2} \eta(\theta_1,\theta_2)\mathbf{p}(\cdot,\theta_1,\theta_2) = \sum_{\theta_1\in\Theta_1}\left(\sum_{\theta_2\in\Theta_2}\eta(\theta_1,\theta_2)\right)\hat{\mathbf{p}}(\cdot,\theta_1) = \hat{\mathbf{p}}_{\eta_1}$$

$$\mathbf{q}_\eta = \sum_{(\theta_1,\theta_2)\in\Theta_1\times\Theta_2} \eta(\theta_1,\theta_2)\mathbf{q}(\cdot,\theta_1,\theta_2) = \sum_{\theta_2\in\Theta_2}\left(\sum_{\theta_1\in\Theta_1}\eta(\theta_1,\theta_2)\right)\hat{\mathbf{q}}(\cdot,\theta_2) = \hat{\mathbf{q}}_{\eta_2}$$

For the if direction, suppose $\mathbf{p}$ does not prior-by-prior dominates $\mathbf{q}$, then there exists $\eta \in \Delta(\Theta_1 \times \Theta_2)$ such that $\mathbf{p}_\eta$ is not Blackwell more informative than $\mathbf{q}_\eta$. This gives us a pair of marginal distributions $(\eta_1,\eta_2) \in \Delta(\Theta_1) \times \Delta(\Theta_2)$ such that $\hat{\mathbf{p}}_{\eta_1}$ is not Blackwell more informative than $\hat{\mathbf{q}}_{\eta_2}$. Taking the contrapositive completes the proof.

For the only if direction, suppose there exists some pair of marginal distributions $(\mu,\nu) \in \Delta(\Theta_1) \times \Delta(\Theta_2)$ such that $\hat{\mathbf{p}}_\mu$ is not Blackwell more informative than $\hat{\mathbf{q}}_\nu$, then we can construct a joint distribution η that induces (μ,ν) such that $\mathbf{p}_\eta$ is not Blackwell more informative than $\mathbf{q}_\eta$, making it impossible for $\mathbf{p}$ to prior-by-prior dominate $\mathbf{q}$. Taking the contrapositive completes the proof. □

B.2.3 Proof of Theorem 2.6

Proof. To prove that $1 \implies 2$, fix any (A, u, π),

$$\begin{aligned}
\max_{\sigma\in A_S}\min_{p\in P} U(\sigma,p) &= \max_{\sigma\in A_S}\min_{p\in \mathrm{conv}(P)} U(\sigma,p) \\
&= \min_{p\in \mathrm{conv}(P)}\max_{\sigma\in A_S} U(\sigma,p) && \text{(minimax theorem)} \\
&= \max_{\sigma\in A_S} U(\sigma,p_*) && \text{(for some } p_*\in \mathrm{conv}(P)) \\
&\geq \max_{\sigma'\in A_{S'}} U(\sigma',p'_*) && \text{(for some } p'_*\in \mathrm{conv}(P') \text{ by condition 2)} \\
&\geq \min_{p'\in \mathrm{conv}(P')}\max_{\sigma'\in A_{S'}} U(\sigma',p') \\
&= \max_{\sigma'\in A_{S'}}\min_{p'\in \mathrm{conv}(P')} U(\sigma',p') = \max_{\sigma'\in A_{S'}}\min_{p'\in P'} U(\sigma',p')
\end{aligned}$$

where the second equality follows from von Neumann's minimax theorem since both $\mathrm{conv}(P)$ and A_S are compact and convex. This concludes the proof for $1 \implies 2$.

Then we prove $2 \implies 1$ by proving its contrapositive.

Suppose there exists $p_0 \in \mathrm{conv}(P)$ such that p_0 is not Blackwell more informative than any $p' \in \mathrm{conv}(P')$, we want to construct a triplet (A, u, π) in which

$$\max_{\sigma\in A_S}\min_{p\in P} U(\sigma,P) < \max_{\sigma'\in A_{S'}}\min_{p'\in P'} U(\sigma',P').$$

Fix $A = S'$, that is, the action space is just the set of signal realizations for P'. Then the sets of action plans are $A_S = \{\sigma \mid \sigma : S \to \Delta(S')\},\ A_{S'} = \{\sigma' \mid \sigma' : S' \to \Delta(S')\}$. Consider an action plan $r \in A_{S'}$ defined by $r(s'_i \mid s'_j) = \mathbf{1}[i = j]$, that is, r is the action plan that just reports the signal realization. Then for any $p' \in P'$, $r \circ p' = p'$ since

$$(r \circ p')(s'_i \mid \omega) = \sum_{s'\in S'} r(s'_i \mid s')p'(s' \mid \omega) = p'(s'_i \mid \omega),\ \forall s'_i \in A \text{ and } \omega \in \Omega.$$

Let $\Lambda_{P',r} := \{r \circ p' \mid p' \in \text{conv}(P')\}$. That is, $\Lambda_{P',r}$ is the set of all probability measures over A conditional on Ω that can be induced by action plan r together with some Blackwell experiment. Then $\Lambda_{P',r} = \text{conv}(P')$. Similarly, let $\Lambda_{p_0} := \{\sigma \circ p_0 \mid \sigma \in A_S\}$.

For any $u : \Omega \times S' \to \mathbb{R}$ and $\pi \in \Delta(\Omega)$,

$$\begin{aligned}
\max_{\sigma \in A_S} U(\sigma, p_0) &= \max_{\sigma \in A_S} \sum_{s \in S} \sum_{\omega \in \Omega} \pi(\omega) p_0(s \mid \omega) \sum_{a \in A} \sigma(a \mid s) u(\omega, a) \\
&= \max_{\sigma \in A_S} \sum_{\omega \in \Omega} \left(\sum_{a \in S'} u(\omega, a) \underbrace{\sum_{s \in S} \sigma(a \mid s) p_0(s \mid \omega)}_{\sigma \circ p_0} \right) \pi(\omega) \\
&= \max_{\lambda \in \Lambda_{p_0}} \sum_{\omega \in \Omega} \left(\sum_{a \in S'} u(\omega, a) \lambda(a \mid \omega) \right) \pi(\omega)
\end{aligned}$$

where the last equality follows from the one-to-one correspondence of A_S and Λ_{p_0}.

Since p_0 is not more informative than p' for any $p' \in \text{conv}(P')$, $\Lambda_{p_0} \cap \Lambda_{P',r} = \emptyset$. If not, there must exist λ in their intersection $\Lambda_{p_0} \cap \Lambda_{P',r}$, which further indicates the existence of an action plan $\sigma : S \to \Delta(S')$ and an experiment $p' \in \text{conv}(P')$ such that $\sigma \circ p_0 = \lambda = p'$, that is, some $p' \in \text{conv}(P')$ is a garbling of p_0. Contradiction.

But both Λ_{p_0} and $\Lambda_{P',r}$ are compact and convex subsets of $\mathbb{R}^{|\Omega| \times |S'|}$, hence we can apply the separating hyperplane theorem and conclude that there exists a nonzero vector $v \in \mathbb{R}^{|\Omega| \times |S'|}$ and real numbers $c_1 < c_2$ such that

$$\sum_{\omega \in \Omega} \sum_{a \in S'} v(\omega, a) \lambda(a \mid \omega) < c_1, \ \sum_{\omega \in \Omega} \sum_{a \in S'} v(\omega, a) \lambda'(a \mid \omega) > c_2, \ \forall \lambda \in \Lambda_{p_0}, \ \forall \lambda' \in \Lambda_{P',r}.$$

Consider $A = S'$, $u = v$ as given above, and $\pi = \text{uniform}(\Omega)$.

$$\max_{\sigma \in A_S} U(\sigma, p_0) = \max_{\lambda \in \Lambda_{p_0}} \sum_{\omega \in \Omega} \left(\sum_{a \in S'} v(\omega, a) \lambda(a \mid \omega) \right) \pi(\omega)$$

$$= \frac{1}{|\Omega|} \cdot \max_{\lambda \in \Lambda_{p_0}} \sum_{\omega \in \Omega} \sum_{a \in S'} v(\omega, a) \lambda(a \mid \omega)$$

$$< \frac{1}{|\Omega|} \cdot \min_{\lambda' \in \Lambda_{P',r}} \sum_{\omega \in \Omega} \sum_{a \in S'} v(\omega, a) \lambda'(a \mid \omega) = \min_{p' \in \text{conv}(P')} U(r, p')$$

where the inequality follows from the separation result above and the last equality follows from the one-to-one correspondence of $\Lambda_{P',r}$ and $\text{conv}(P')$.

Therefore, with $(A, u, \pi) = (S', v, \text{uniform})$,

$$\max_{\sigma \in A_S} \min_{p \in P} U(\sigma, p) = \max_{\sigma \in A_S} \min_{p \in \text{conv}(P)} U(\sigma, p)$$

$$= \min_{p \in \text{conv}(P)} \max_{\sigma \in A_S} U(\sigma, p)$$

$$\leq \max_{\sigma \in A_S} U(\sigma, p_0)$$

$$< \min_{p' \in \text{conv}(P')} U(r, p') = \min_{p' \in P'} U(r, p') \leq \max_{\sigma' \in A_{S'}} \min_{p' \in P'} U(\sigma'; P')$$

where the first and the last equalties follow from the facts that $U(\cdot, \cdot)$ is linear in its second argument. The second inequality follows from von Neumann's minimax theorem since both $\text{conv}(P)$ and A_S are convex and compact and $U(\cdot, \cdot)$ is linear in both arguments. This completes the proof of the contrapositive. □

B.2.4 Proof of Proposition 2.7

Proof of Proposition 2.7. We focus on $\mathbf{p}$ and P, the proof for $\mathbf{p}'$ and P' is the same. Fix A, u, π and σ. Under the construction in (2.14),

$$V_W \left(\left\{ U\big(\sigma, \mathbf{p}(\cdot, \theta)\big) \right\}_{\theta \in \Theta} \right) = \min_{\theta \in \Theta} U\big(\sigma, \mathbf{p}(\cdot, \theta)\big)$$

$$= \min_{(p, p') \in P \times P'} U\big(\sigma, \mathbf{p}(\cdot, (p, p'))\big) = \min_{p \in P} U(\sigma, p).$$

Under the construction in (2.15),

$$\min_{p\in P} U(\sigma,p)=\min_{\theta\in\Theta} U\big(\sigma,\mathbf{p}(\cdot,\theta)\big)=V_W\left(\big\{U\big(\sigma,\mathbf{p}(\cdot,\theta)\big)\big\}_{\theta\in\Theta}\right).$$

This completes the proof. □

B.2.5 Proof or Proposition 2.9

Proof. To see $1 \implies 2$: let σ' be any action plan made for $\mathbf{p}'$ and let γ be the garbling that transforms $\mathbf{p}$ to $\mathbf{p}'$. Consider $\sigma'\circ\gamma$, $\sigma'\circ\gamma$ is a feasible action plan for $\mathbf{p}$, moreover, for any $(\theta,\omega)\in\Theta\times\Omega$,

$$\begin{aligned}
\hat{U}(\sigma'\circ\gamma,\mathbf{p}(\cdot,\theta),\omega) &= \sum_{s\in S}\sum_{a\in A}\mathbf{p}(s\mid\omega,\theta)(\sigma'\circ\gamma)(a\mid s)u(\omega,a)\\
&= \sum_{s\in S}\sum_{a\in A}\mathbf{p}(s\mid\omega,\theta)\sum_{s'\in S'}\sigma'(a\mid s')\gamma(s'\mid s)u(\omega,a)\\
&= \sum_{s'\in S'}\sum_{a\in A}\left(\sum_{s\in S}\mathbf{p}(s\mid\omega,\theta)\gamma(s'\mid s)\right)\sigma'(a\mid s')u(\omega,a)\\
&= \sum_{s'\in S'}\sum_{a\in A}\mathbf{p}'(s'\mid\omega,\theta)\sigma'(a\mid s')u(\omega,a)=\hat{U}(\sigma',\mathbf{p}'(\cdot,\theta),\omega)
\end{aligned}$$

That is, by taking the action plan $\sigma'\circ\gamma$, the DM obtains the same conditional expected utilities in every pair of states (ω,θ). This together with the assumption that the aggregator $\hat{V}$ is monotone completes the proof.

The equivalence of 2 and 3 is straightforward.

To see $3 \implies 4$: Comparing to our model in Section 3, the combination of a unique prior π over Ω and an aggregator V over $\mathbb{R}^{\Theta}$ is just one special case of the more general aggregator $\hat{V}$ over $\mathbb{R}^{\Omega\times\Theta}$. Therefore, the set of possible (A,u,π,V) is expanded, and prior-by-prior dominance must still be necessary for guaranteeing higher ex-ante utility. □

B.2.6 Proof of Proposition 2.10

Proof. Fix any finite set of actions $A = \{a_1, a_2, \dots, a_n\}$, and let the state-dependent utility $u : \Omega \times A \to \mathbb{R}$ be summarized by

$$u(\omega_i, a_j) = u_{ij} \text{ for } i \in \{1, 2\} \text{ and } j \in \{1, \dots, n\}.$$

Let σ denote the DM's action plan facing $\mathbf{p}$, with $\sigma_{ij} := \sigma(a_j \mid s_i)$. That is, σ_{ij} is the probability that the DM plays a_j after observing signal s_i. Note that $\sum_j \sigma_{ij} = 1$ for $i \in \{1, 2\}$. Let σ' denote the DM's action plan facing $\mathbf{p}'$, with σ'_{ij} defined accordingly. Write $\hat{U}^{\sigma}_{\mathbf{p}}(\theta, \omega)$ to get a more compact notation for conditional expected utilities , that is,

$$\hat{U}^{\sigma}_{\mathbf{p}}(\theta, \omega) := \hat{U}\big(\sigma, \mathbf{p}(\cdot, \theta), \omega\big).$$

Then we have

$$\hat{U}^{\sigma}_{\mathbf{p}}(\theta, \omega_j) = \sum_{k=1}^{n} \sigma_{jk} u_{jk}, \ \forall \theta \in \{\theta_1, \theta_2\}, \ j \in \{1, 2\}$$

$$\hat{U}^{\sigma'}_{\mathbf{p}'}(\theta_1, \omega_j) = \sum_{k=1}^{n} \sigma'_{jk} u_{jk}, \ \forall j \in \{1, 2\}$$

$$\hat{U}^{\sigma'}_{\mathbf{p}'}(\theta_2, \omega_1) = \sum_{k=1}^{n} \sigma'_{2k} u_{1k} \quad \text{and} \quad \hat{U}^{\sigma'}_{\mathbf{p}'}(\theta_2, \omega_2) = \sum_{k=1}^{n} \sigma'_{1k} u_{2k}$$

Thus, for any $\sigma' \in A_{S'}$, we can find $\sigma \in A_S$ such that

$$\hat{U}^{\sigma}_{\mathbf{p}}(\theta, \omega) \geq \hat{U}^{\sigma'}_{\mathbf{p}'}(\theta, \omega), \text{ for all } (\theta, \omega) \in \Theta \times \Omega.$$

To achieve this, let $\bar{u}_i := \max_{j \in \{1,\dots,n\}} u_{ij}$ and $n^*_i \in \arg\max_{j \in \{1,\dots,n\}} u_{ij}$, and consider $\sigma^* \in A_S$ be defined by $\sigma^*_{ij} = \mathbf{1}[j = n^*_i]$. Let σ' be an arbitrary action plan in $A_{S'}$, then

$$\hat{U}^{\sigma^*}_{\mathbf{p}}(\theta_1, \omega_1) = \bar{u}_1 = \sum_{k=1}^{n} \sigma'_{1k} \bar{u}_1 \geq \sum_{k=1}^{n} \sigma'_{1k} u_{1k} = \hat{U}^{\sigma'}_{\mathbf{p}'}(\theta_1, \omega_1)$$

$$\hat{U}^{\sigma^*}_{\mathbf{p}}(\theta_1, \omega_2) = \bar{u}_2 = \sum_{k=1}^{n} \sigma'_{2k} \bar{u}_2 \geq \sum_{k=1}^{n} \sigma'_{2k} u_{2k} = \hat{U}^{\sigma'}_{\mathbf{p}'}(\theta_1, \omega_2)$$

$$\hat{U}^{\sigma^*}_{\mathbf{p}}(\theta_2, \omega_1) = \bar{u}_1 = \sum_{k=1}^{n} \sigma'_{2k} \bar{u}_1 \geq \sum_{k=1}^{n} \sigma'_{2k} u_{1k} = \hat{U}^{\sigma'}_{\mathbf{p}'}(\theta_2, \omega_1)$$

$$\hat{U}^{\sigma^*}_{\mathbf{p}}(\theta_2, \omega_2) = \bar{u}_2 = \sum_{k=1}^{n} \sigma'_{1k} \bar{u}_2 \geq \sum_{k=1}^{n} \sigma'_{1k} u_{2k} = \hat{U}^{\sigma'}_{\mathbf{p}'}(\theta_2, \omega_2)$$

Then for any monotone aggregator $\hat{V} : \mathbb{R}^{\Omega \times \Theta} \to \mathbb{R}$,

$$\sup_{\sigma \in A_S} \hat{V}\left(\hat{U}^{\sigma}_{\mathbf{p}}(\cdot, \cdot)\right) \geq \hat{V}\left(\hat{U}^{\sigma^*}_{\mathbf{p}}(\cdot, \cdot)\right) \geq \sup_{\sigma' \in A_{S'}} \hat{V}\left(\hat{U}^{\sigma'}_{\mathbf{p}'}(\cdot, \cdot)\right).$$

That is, $\mathbf{p}$ gives weakly higher ex-ante utility than $\mathbf{p}'$ for any possible $(A, u, \hat{V})$. □

B.2.7 Proof of Theorem B.1

We prove Theorem B.1 by showing that $1 \Longrightarrow 2 \Longrightarrow 3 \Longrightarrow 1$.

Proof. Recall that Δ_0 is the set of all probability measures over Θ with finite support and $\mathbf{p} \trianglerighteq_{PBP} \mathbf{p}'$ if $\mathbf{p}_\mu$ is Blackwell more informative than $\mathbf{p}'_\mu$ for all $\mu \in \Delta_0$.

To see $1 \Longrightarrow 2$:

Let Γ denote the set of all garblings, that is, $\Gamma := \{\gamma \mid \gamma : S \to \Delta(S')\}$. Γ is compact and convex since both S and S' are finite. Fix A, u, π and $\sigma' \in A_{S'}$, we define an auxiliary function $H : \Gamma \times \Delta_0 \to \mathbb{R}$ by

$$H(\gamma, \mu) := U(\sigma' \circ \gamma, \mathbf{p}_\mu) - U(\sigma', \mathbf{p}'_\mu).$$

That is, $H(\gamma, \mu)$ is difference in the expected utilities of $\mathbf{p}$ and $\mathbf{p}'$ conditional on belief $\mu \in \Theta$ and garbling γ being applied to the action plan made for $\mathbf{p}'$.

To prove 1 implies 2, it suffices to show that

$$\max_{\gamma \in \Gamma} \inf_{\mu \in \Delta_0} H(\gamma, \mu) \geq 0.$$

To show that the left hand side is well defined and nonnegative, we invoke the Kneser-Fan minimax theorem for concave-convex functions (Terkelsen, 1972).

Γ is a compact and convex subset of $\mathbb{R}^{|S| \times |S'|}$ under the Euclidean topology. Δ_0 is a convex subset of the vector space $\mathbb{R}^{\Theta}$. For each $\gamma \in \Gamma$, the function $\mu \mapsto -H(\gamma, \mu)$ is linear (hence concave) on Δ_0. For each $\mu \in \Delta_0$, the function $\gamma \mapsto -H(\gamma, \mu)$ is linear (hence convex and continuous since Γ is finite dimensional) on Γ. Therefore, by the Kneser-Fan minimax theorem for concave-convex functions (Terkelsen, 1972, page 411, Corollary 2),

$$\min_{\gamma \in \Gamma} \sup_{\mu \in \Delta_0} -H(\gamma, \mu) = \sup_{\mu \in \Delta_0} \min_{\gamma \in \Gamma} -H(\gamma, \mu),$$

which further indicates that

$$\max_{\gamma \in \Gamma} \inf_{\mu \in \Delta_0} H(\gamma, \mu) = \inf_{\mu \in \Delta_0} \max_{\gamma \in \Gamma} H(\gamma, \mu).$$

But the right hand side of the equation above is non-negative, since for any $\mu \in \Delta_0$, the prior-by-prior dominance condition guarantees that there exists some $\gamma_\mu \in \Gamma$ such that $\mathbf{p}'_\mu = \gamma_\mu \circ \mathbf{p}_\mu$ and this garbling γ_μ guarantees that $H(\gamma_\mu, \mu) = 0$. This completes the proof that $1 \Longrightarrow 2$ since $\sigma' \circ \gamma \in A_S$ is a valid action plan for $\mathbf{p}$.

To see $2 \Longrightarrow 3$**:**

Fix any (A, u, π, V) with $V \in \mathcal{V}_{Mono}$ and let v^* denote the ex-ante utility for $\mathbf{p}'$, that is,

$$v^* := \sup_{\sigma' \in A_{S'}} V\left(\left\{U\big(\sigma', \mathbf{p}'(\cdot, \theta)\big)\right\}_{\theta \in \Theta}\right).$$

Then for any $\varepsilon > 0$, there exists $\sigma'_\varepsilon \in A_{S'}$ such that

$$V\left(\left\{U\big(\sigma'_\varepsilon, \mathbf{p}'(\cdot, \theta)\big)\right\}_{\theta \in \Theta}\right) > v^* - \varepsilon.$$

By condition 2 in Theorem B.1, there exists an action plan $\sigma_\varepsilon \in A_S$ such that

$$U\big(\sigma_\varepsilon, \mathbf{p}(\cdot,\theta)\big) \geq U\big(\sigma'_\varepsilon, \mathbf{p}'(\cdot,\theta)\big),\ \forall \theta \in \Theta.$$

By the monotonicity of V, this further indicates that

$$V\left(\big\{U\big(\sigma_\varepsilon, \mathbf{p}(\cdot,\theta)\big)\big\}_{\theta\in\Theta}\right) \geq V\left(\big\{U\big(\sigma'_\varepsilon, \mathbf{p}'(\cdot,\theta)\big)\big\}_{\theta\in\Theta}\right) > v^* - \varepsilon.$$

Since this holds for any $\varepsilon > 0$, we have $\sup_{\sigma\in A_S} V\left(\big\{U\big(\sigma, \mathbf{p}(\cdot,\theta)\big)\big\}_{\theta\in\Theta}\right) \geq v^*$.

To see $3 \Longrightarrow 1$:

It suffices to prove the necessity of prior-by-prior dominance when $\mathcal{V} = \mathcal{V}_0$. That is, if $\mathbf{p}$ is preferred to $\mathbf{p}'$ by every decision maker with $V \in \mathcal{V}_0$, then $\mathbf{p} \trianglerighteq_{PBP} \mathbf{p}'$.

Suppose by contradiction that it is not the case $\mathbf{p} \trianglerighteq_{PBP} \mathbf{p}'$, then there must exist some $\mu \in \Delta_0$ such that $\mathbf{p}_\mu$ is not Blackwell more informative than $\mathbf{p}'_\mu$. Fixing this belief μ and its corresponding aggregator $V_\mu \in \mathcal{V}_0$, that is, the DM believes that μ is the correct distribution over Θ and use aggregator V_μ to evaluate action plans. Thus, this DM's ex-ante utility is

$$\begin{aligned}\sup_{\sigma\in A_S}\ V_\mu\left(\big\{U\big(\sigma, \mathbf{p}(\cdot,\theta)\big)\big\}_{\theta\in\Theta}\right) &= \sup_{\sigma\in A_S} \sum_{\theta\in\Theta} U\big(\sigma, \mathbf{p}(\cdot,\theta)\big)\mu(\theta)\\ &= \sup_{\sigma\in A_S}\ U\left(\sigma, \sum_{\theta\in\Theta} \mathbf{p}(\cdot,\theta)\mu(\theta)\right)\\ &= \sup_{\sigma\in A_S}\ U(\sigma, \mathbf{p}_\mu) = \max_{\sigma\in A_S}\ U(\sigma, \mathbf{p}_\mu)\end{aligned}$$

Since $\mathbf{p}_\mu$ is not Blackwell more informative than $\mathbf{p}'_\mu$, then by Theorem 2.1, there must exist a triplet (A, u, π) such that

$$\max_{\sigma\in A_S}\ U(\sigma, \mathbf{p}_\mu) < \max_{\sigma'\in A_{S'}}\ U(\sigma', \mathbf{p}'_\mu),$$

which further indicates that in (A, u, π, V_μ),

$$\sup_{\sigma\in A_S}\ V_\mu\left(\big\{U\big(\sigma, \mathbf{p}(\cdot,\theta)\big)\big\}_{\theta\in\Theta}\right) < \sup_{\sigma'\in A_{S'}}\ V_\mu\left(\big\{U\big(\sigma', \mathbf{p}'(\cdot,\theta)\big)\big\}_{\theta\in\Theta}\right)$$

which is a direct contradiction to our assumption that $\mathbf{p}$ gives weakly higher ex-ante utility for any (A, u, π, V) with $V \in \mathcal{V}_0$. This completes the proof that prior-by-prior dominance is necessary for any class of aggregators $\mathcal{V} \supset \mathcal{V}_0$. $\square$

Appendix C

Appendix to Chapter 3

C.1 Proof of Proposition 3.1

To prove Proposition 3.1, it is useful to first establish the following result: if two distinct messages lead to the same action of the Receiver in some feasible information structure, then a new information structure, created by relabeling those messages as one and adding up their conditional probabilities as the new conditional probability, is also feasible and gives the same persuasion payoff to the Sender.

Formally, fix a constrained persuasion problem $\mathcal{P}$ and a feasible information structure (M, π), with

$$M = \{m_1, \ldots, m_n\},\ n \geq 3;$$
$$\pi : \Theta \to \Delta M \text{ with } \pi(m_i \mid \theta) \in [\alpha, 1-\alpha],\ \forall i \text{ and } \forall \theta.$$

If two distinct messages $m_i, m_j \in M$ leads to the same action a of the Receiver, we can define another information structure (M', π') by

$$M' = \{m_1, \ldots, \widehat{m_i}, \ldots, \widehat{m_j}, \ldots, m_n, a\}$$
$$\pi' : \Theta \to \Delta M' \text{ with } \pi'(a \mid \theta) = \pi(m_i \mid \theta) + \pi(m_j \mid \theta),\ \forall \theta \in \Theta$$
$$\text{and } \pi'(m_k \mid \theta) = \pi(m_k \mid \theta),\ \forall k \notin \{i, j\} \text{ and } \forall \theta \in \Theta$$

Lemma C.1. *(M', π') is feasible and achieves the same persuasion payoff as (M, π).*

Proof of Lemma C.1. We first prove the feasibility of (M', π'). For $k \notin \{i, j\}$,

$$\pi'(m_k \mid \theta) = \pi(m_k \mid \theta) \in [\alpha, 1 - \alpha], \ \forall \theta \in \Theta,$$

where the containment follows from (M, π) being a feasible information structure.

For $a \in M'$, it is clear that

$$\pi'(a \mid \theta) = \pi(m_i \mid \theta) + \pi(m_j \mid \theta) \geq 2\alpha > \alpha, \ \forall \theta \in \Theta.$$

To show $\pi'(a \mid \theta) \leq 1 - \alpha$ for all θ, recall that we have the assumption that $|M| \geq 3$, hence $|M'| \geq 2$, and there exists at least another different message $m \neq a$. Therefore,

$$\pi'(a \mid \theta) \leq 1 - \pi'(m \mid \theta) \leq 1 - \alpha, \ \forall \theta \in \Theta,$$

where the first inequality follows from $\pi'(\cdot \mid \theta)$ being a probability distribution for each θ and the second inequality follows from the fact that $\pi'(m \mid \theta) = \pi(m \mid \theta)$ and π is feasible. This completes the proof that (M', π') as constructed above is feasible.

We then show that (M', π') achieves the same persuasion payoff as (M, π).

Step 1. Show the message $a \in M'$ leads to action a.

Let a_i^* denote the action message m_i leads to, i.e.,

$$a_i^* \in \hat{a}(\mu_{m_i}) := \arg\max_{a \in A} \sum_{\theta \in \Theta} \mu_{m_i}(\theta) u_R(a, \theta), \quad \text{and}$$

$$\sum_{\theta \in \Theta} \mu_{m_i}(\theta) u_S(a_i^*, \theta) > \sum_{\theta \in \Theta} \mu_{m_i}(\theta) u_S(a, \theta), \ \forall a \in \hat{a}(\mu_{m_i}) \setminus \{a_i^*\}$$

a_i^* is the unique Sender-preferred Receiver-optimal action under belief μ_{m_i} after observing message m_i.

By assumption, $a_i^* = a_j^* = a$. Under the newly constructed information structure (M', π'), if message a is sent, the Bayes' rule gives the posterior μ_a as

$$\mu_a(\theta) = \frac{\mu(\theta) \pi'(a \mid \theta)}{\sum_{\theta' \in \Theta} \mu(\theta') \pi'(a \mid \theta')} = \frac{\mu(\theta)[\pi(m_i \mid \theta) + \pi(m_j \mid \theta)]}{\sum_{\theta' \in \Theta} \mu(\theta')[\pi(m_i \mid \theta') + \pi(m_j \mid \theta')]}.$$

It should then be clear that $a \in \hat{a}(\mu_a)$, that is, a is indeed a best response of the Receiver when he has posterior belief μ_a, since for any $a' \in A$

$$\begin{aligned}
\sum_{\theta\in\Theta} \mu_a(\theta)u_R(a,\theta) &= \frac{\sum_{\theta\in\Theta}\mu(\theta)\pi(m_i \mid \theta)u_R(a,\theta) + \sum_{\theta\in\Theta}\mu(\theta)\pi(m_j \mid \theta)u_R(a,\theta)}{\sum_{\theta\in\Theta}\mu(\theta)[\pi(m_i \mid \theta) + \pi(m_j \mid \theta)]} \\
&\geq \frac{\sum_{\theta\in\Theta}\mu(\theta)\pi(m_i \mid \theta)u_R(a',\theta) + \sum_{\theta\in\Theta}\mu(\theta)\pi(m_j \mid \theta)u_R(a',\theta)}{\sum_{\theta\in\Theta}\mu(\theta)[\pi(m_i \mid \theta) + \pi(m_j \mid \theta)]} \\
&= \sum_{\theta\in\Theta} \mu_a(\theta)u_R(a',\theta)
\end{aligned}$$

where the inequality follows from the fact that $a = a_i^* \in \hat{a}(\mu_{m_i})$ and

$$\begin{aligned}
&a \in \hat{a}(\mu_{m_i}) \\
\iff &\sum_{\theta\in\Theta}\mu(\theta)\pi(m_i \mid \theta)u_R(a,\theta) \geq \sum_{\theta\in\Theta}\mu(\theta)\pi(m_i \mid \theta)u_R(a',\theta),\ \forall a' \in A.
\end{aligned}$$

This part of the proof largely resembles the proof of Proposition 1 in KG. The result can be written more compactly as

$$\hat{a}(\mu_a) \supseteq \hat{a}(\mu_{m_i}) \cap \hat{a}(\mu_{m_j}).$$

Lastly, we need to actually check that a is the unique Sender-preferred Receiver-optimal action under belief μ_a. That is, we need to verify that

$$a' \in \hat{a}(\mu_a) \setminus \{a\} \implies \sum_{\theta\in\Theta}\mu_a(\theta)u_S(a,\theta) > \sum_{\theta\in\Theta}\mu_a(\theta)u_S(a',\theta).$$

We show this by showing that

$$\hat{a}(\mu_a) \subseteq \hat{a}(\mu_{m_i}) \cap \hat{a}(\mu_{m_j}).$$

Suppose $a' \in \hat{a}(\mu_a)$ and $a' \neq a$, want to show: $a' \in \hat{a}(\mu_{m_i}) \cap \hat{a}(\mu_{m_j})$.

Both a and a' are in the set $\hat{a}(\mu_a)$ implies that

$$\sum_{\theta\in\Theta}\mu_a(\theta)u_R(a,\theta) = \sum_{\theta\in\Theta}\mu_a(\theta)u_R(a',\theta).$$

Then from the previous derivation

$$\begin{aligned}&\sum_{\theta\in\Theta}\mu(\theta)\pi(m_i \mid \theta)u_R(a,\theta)+\sum_{\theta\in\Theta}\mu(\theta)\pi(m_j \mid \theta)u_R(a,\theta)\\ =&\sum_{\theta\in\Theta}\mu(\theta)\pi(m_i \mid \theta)u_R(a',\theta)+\sum_{\theta\in\Theta}\mu(\theta)\pi(m_j \mid \theta)u_R(a',\theta).\end{aligned} \tag{$\star$}$$

Suppose by contradiction that $a' \notin \hat{a}(\mu_{m_i})$, then $a \in \hat{a}(\mu_{m_i})$ implies that

$$\sum_{\theta\in\Theta}\mu(\theta)\pi(m_i \mid \theta)u_R(a',\theta) < \sum_{\theta\in\Theta}\mu(\theta)\pi(m_i \mid \theta)u_R(a,\theta).$$

Moreover, $a \in \hat{a}(\mu_{m_j})$ implies that

$$\sum_{\theta\in\Theta}\mu(\theta)\pi(m_j \mid \theta)u_R(a',\theta) \leq \sum_{\theta\in\Theta}\mu(\theta)\pi(m_j \mid \theta)u_R(a,\theta).$$

The combination of the last two inequalities constitute a direct contradiction to equation ($\star$). Hence the hypothesis cannot be true and $a' \in \hat{a}(\mu_{m_i})$. Similar arguments guarantee that $a' \in \hat{a}(\mu_{m_j})$. Hence $a' \in \hat{a}(\mu_{m_i}) \cap \hat{a}(\mu_{m_j})$. This completes the proof that

$$\hat{a}(\mu_a) \subseteq \hat{a}(\mu_{m_i}) \cap \hat{a}(\mu_{m_j}).$$

That is, every action that is Receiver-optimal under belief μ_a is Receiver-optimal under belief μ_{m_i} and μ_{m_j}. But as assumed, both m_i and m_j lead to action a. Hence

$$\sum_{\theta\in\Theta}\mu_a(\theta)u_s(a,\theta) > \sum_{\theta\in\Theta}\mu_a(\theta),\ \forall a \in \hat{a}\mu_a \setminus \{a\}.$$

This completes the proof since it is assumed that a is the unique Sender-preferred Receiver-optimal action after observing m_i or m_j.[1]

[1] Note that the uniqueness assumption on the Sender-preferred Receiver-optimal action is only used in the last step of the proof. Hence the same proof will go through as long as the Receiver's tie-breaking rule (choice function) is rationalizable.

Step 2. Show that (M', π') achieves the same persuasion payoff as (M, π).

This follows from applying the result from step 1,

$$
\begin{aligned}
v_{M,\pi} := & \sum_{m \in M} \left[\underbrace{\sum_{\theta \in \Theta} \mu(\theta)\pi(m \mid \theta)}_{=\mathsf{P}(\text{message } m \text{ is sent})} \left(\sum_{\theta \in \Theta} \mu_m(\theta) u_S(a^*(m), \theta) \right) \right] \\
= & \sum_{m \in M \setminus \{m_i, m_j\}} \left[\sum_{\theta \in \Theta} \mu(\theta)\pi(m \mid \theta) \left(\sum_{\theta \in \Theta} \mu_m(\theta) u_S(a^*(m), \theta) \right) \right] \\
& + \underbrace{\sum_{\theta \in \Theta} \left[\mu(\theta)\pi(m_i \mid \theta) \left(\sum_{\theta \in \Theta} \mu_{m_i}(\theta) u_S(a, \theta) \right) \right]}_{=\sum_{\theta \in \Theta} \mu(\theta)\pi(m_i \mid \theta) u_S(a,\theta)} \\
& + \underbrace{\sum_{\theta \in \Theta} \left[\mu(\theta)\pi(m_j \mid \theta) \left(\sum_{\theta \in \Theta} \mu_{m_j}(\theta) u_S(a, \theta) \right) \right]}_{=\sum_{\theta \in \Theta} \mu(\theta)\pi(m_j \mid \theta) u_S(a,\theta)} \\
= & \sum_{m \in M' \setminus \{a\}} \left[\sum_{\theta \in \Theta} \mu(\theta)\pi'(m \mid \theta) \left(\sum_{\theta \in \Theta} \mu_m(\theta) u_S(a^*(m), \theta) \right) \right] \\
& + \sum_{\theta \in \Theta} \mu(\theta) \underbrace{\left[\pi(m_i \mid \theta) + \pi(m_j \mid \theta) \right]}_{=\pi'(a \mid \theta)} u_S(a, \theta) \\
= & \sum_{m \in M' \setminus \{a\}} \left[\sum_{\theta \in \Theta} \mu(\theta)\pi'(m \mid \theta) \left(\sum_{\theta \in \Theta} \mu_m(\theta) u_S(a^*(m), \theta) \right) \right] \\
& + \sum_{\theta \in \Theta} \left[\mu(\theta)\pi'(a \mid \theta) \left(\sum_{\theta \in \Theta} \mu_a(\theta) u_S(a, \theta) \right) \right] \\
= & \sum_{m \in M'} \left[\sum_{\theta \in \Theta} \mu(\theta)\pi'(m \mid \theta) \left(\sum_{\theta \in \Theta} \mu_m(\theta) u_S(a^*(m), \theta) \right) \right] \\
= & v_{M',\pi'}
\end{aligned}
$$

This completes Step 2 and the proof of Lemma C.1. □

Proposition 3.1 follows from applying Lemma C.1 repeatedly to the original information structure (M', π') with $|M'| > |A|$.

Proof of Proposition 3.1. By assumption, (M', π') is feasible and achieves persuasion payoff v. Since $|M'| > |A|$, there exists a pair of distinct messages $m'_1, m'_2 \in M'$ that leads to the same action (otherwise we should have $|M'| \leq |A|$), apply Lemma C.1 to get a new information structure (M'', π'') with $|M''| = |M'| - 1$. The same arguments go through for any intermediate information structure with $|M| > |A|$, and applying Lemma C.1 repeatedly until we reach the final information structure with $|M| = |A|$. □

An alternative but similar proof involves relabeling not only two, but all distinct messages that lead to the same action by that action and add up their conditional probabilities. The motivation we choose the current method is that we want to end with $|M| = |A|$, whereas the alternative proof might end up with some $|M| < |A|$ when only a proper subset of actions are lead to by the original information structure.

C.2 Proof of Proposition 3.2

Proof of Proposition 3.2.

Fix some α (parameter for the constraint) and μ (prior belief that θ_1 is the true state), with the previous calculations, this gives us the bounds on feasible posterior beliefs,

$$\mathrm{LB}_\alpha(\mu) = \frac{\alpha\mu}{1 - \alpha + (2\alpha - 1)\mu}, \quad \mathrm{UB}_\alpha(\mu) = \frac{(1-\alpha)\mu}{\alpha + (1 - 2\alpha)\mu}.$$

Then by the similar arguments we used above, fix any $\mu_1 \in [\mathrm{LB}_\alpha(\mu), \mathrm{UB}_\alpha(\mu)]$, there exist $p, q \in [\alpha, 1 - \alpha]$ such that $\mu_1 = \frac{p\mu}{p\mu + q(1-\mu)}$, i.e., posterior μ_1 is feasible.

$$\mu_1 = \frac{p\mu}{p\mu + q(1-\mu)} \iff p(1-\mu_1)\mu = q\mu_1(1-\mu) \iff q = \left(\frac{1-\mu_1}{\mu_1} \cdot \frac{\mu}{1-\mu}\right) \cdot p.$$

Let $s := \frac{1-\mu_1}{\mu_1} \cdot \frac{\mu}{1-\mu}$. The other belief that can be reached is

$$\mu_2 = \frac{(1-p)\mu}{(1-p)\mu + (1-q)(1-\mu)} = \frac{(1-p)\mu}{(1-p)\mu + (1-sp)(1-\mu)}.$$

First order condition gives

$$\frac{\partial \mu_2}{\partial p} = \frac{-\mu \cdot [(1-p)\mu + (1-sp)(1-\mu)] - (1-p)\mu \cdot [-\mu - s(1-\mu)]}{[(1-p)\mu + (1-sp)(1-\mu)]^2}$$

$$= \cdots = \frac{\mu(1-\mu)(s-1)}{[(1-p)\mu + (1-sp)(1-\mu)]^2}$$

For $\mu_1 \in [\mathrm{LB}_\alpha(\mu), \mu)$, $s = \frac{1-\mu_1}{\mu_1} \cdot \frac{\mu}{1-\mu} > 1$, hence $\partial\mu_2/\partial p > 0$, and μ_2 is strictly increasing in p. Then to get the range of possible μ_2, it suffices to get the feasible range for p. Since $s > 1$, and $p, q = sp \in [\alpha, 1-\alpha]$, the bounds on p and q are

$$\underline{p} = \alpha,\ \underline{q} = s\alpha,\ \overline{p} = \frac{1-\alpha}{s},\ \overline{q} = 1-\alpha.$$

Hence the bounds on μ_2 for a fixed μ_1 are

$$\text{Lower bound} = \frac{(1-\underline{p})\mu}{(1-\underline{p})\mu + (1-s\underline{p})(1-\mu)} = \cdots = \mu \cdot \frac{\mu_1 - \alpha\mu_1}{\mu_1 - \alpha\mu}$$

$$\text{Upper bound} = \frac{(1-\overline{p})\mu}{(1-\overline{p})\mu + (1-s\overline{p})(1-\mu)} = \cdots = \frac{\mu - \mu_1 + \mu_1 \cdot \alpha(1-\mu)}{\mu - \mu_1 + \alpha(1-\mu)}$$

For $\mu_1 \in (\mu, \mathrm{UB}_\alpha(\mu)]$, $s = \frac{1-\mu_1}{\mu_1} \cdot \frac{\mu}{1-\mu} < 1$, hence $\partial\mu_2/\partial p < 0$, and μ_2 is strictly decreasing in p. Then to get the range of possible μ_2, it suffices to get the feasible range for p. Since $s < 1$, and $p, q = sp \in [\alpha, 1-\alpha]$, the bounds on p and q are

$$\underline{p} = \frac{\alpha}{s},\ \underline{q} = \alpha,\ \overline{p} = 1-\alpha,\ \overline{q} = s(1-\alpha).$$

Hence the bounds on μ_2 for a fixed μ_1 are

$$\text{Lower bound} = \frac{(1-\overline{p})\mu}{(1-\overline{p})\mu + (1-s\overline{p})(1-\mu)} = \cdots = \mu \cdot \frac{\alpha\mu_1}{\mu_1 - \mu(1-\alpha)}$$

$$\text{Upper bound} = \frac{(1-\underline{p})\mu}{(1-\underline{p})\mu + (1-s\underline{p})(1-\mu)} = \cdots = \frac{\mu(1-\mu_1) - \mu_1 \cdot \alpha(1-\mu)}{1 - \mu_1 - \alpha(1-\mu)}$$

Last but not least, if μ itself is reached with some positive probability as a posterior, then the pair must be (μ, μ) due to they fact that μ_1 and μ_2 need to average to μ. $\square$

Bibliography

Aliprantis, C. D. and K. C. Border (2006): *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer.

Anscombe, F. J. and R. J. Aumann (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199–205.

Asano, T. and H. Kojima (2019): "Consequentialism and Dynamic Consistency in Updating Ambiguous Beliefs," *Economic Theory*, 68, 223–250.

Aumann, R. J. and M. Maschler (1995): *Repeated Games with Incomplete Information*, MIT press.

Beauchêne, D., J. Li, and M. Li (2019): "Ambiguous Persuasion," *Journal of Economic Theory*, 179, 312–365.

Bell, D. E. (1982): "Regret in Decision Making under Uncertainty," *Operations Research*, 30, 961–981.

Bénabou, R. and J. Tirole (2011): "Identity, Morals, and Taboos: Beliefs as Assets," *The Quarterly Journal of Economics*, 126, 805–855.

Bergemann, D. and S. Morris (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11, 487–522.

Blackwell, D. (1951): "Comparison of Experiments," *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 93–102.

——— (1953): "Equivalent Comparisons of Experiments," *The Annals of Mathematical Statistics*, 265–272.

Bloedel, A. W. and I. R. Segal (2018): "Persuasion with Rational Inattention," *Available at SSRN 3164033*.

Bose, S. and L. Renou (2014): "Mechanism Design with Ambiguous Communication Devices," *Econometrica*, 82, 1853–1872.

Brock, T. C. and J. L. Balloun (1967): "Behavioral Receptivity to Dissonant Information," *Journal of Personality and Social Psychology*, 6, 413.

Brunnermeier, M. K. and J. A. Parker (2005): "Optimal Expectations," *American Economic Review*, 95, 1092–1118.

Buturak, G. and Ö. Evren (2017): "Choice Overload and Asymmetric Regret," *Theoretical Economics*, 12, 1029–1056.

Caplin, A. and J. Leahy (2001): "Psychological Expected Utility Theory and Anticipatory Feelings," *The Quarterly Journal of Economics*, 116, 55–79.

Çelen, B. (2012): "Informativeness of Experiments for MEU," *Journal of Mathematical Economics*, 48, 404–406.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011): "Uncertainty Averse Preferences," *Journal of Economic Theory*, 146, 1275–1330.

Chen, J. Y. (2022): "Biased learning under ambiguous information," *Journal of Economic Theory*, 203, 105492.

Crémer, J. (1982): "A Simple Proof of Blackwell's "Comparison of Experiments" Theorem," *Journal of Economic Theory*, 27, 439–443.

de Oliveira, H. (2018): "Blackwell's Informativeness Theorem Using Diagrams," *Games and Economic Behavior*, 109, 126–131.

Dekel, E., B. L. Lipman, and A. Rustichini (2001): "Representing Preferences with a Unique Subjective State Space," *Econometrica*, 69, 891–934.

——— (2009): "Temptation-Driven Preferences," *The Review of Economic Studies*, 76, 937–971.

Dekel, E., B. L. Lipman, A. Rustichini, and T. Sarver (2007): "Representing Preferences with a Unique Subjective State Space: A Corrigendum," *Econometrica*, 75, 591–600.

Dillenberger, D. (2010): "Preferences for One-Shot Resolution of Uncertainty and Allais-Type Behavior," *Econometrica*, 78, 1973–2004.

Dillenberger, D., J. S. Lleras, P. Sadowski, and N. Takeoka (2014): "A Theory of Subjective Learning," *Journal of Economic Theory*, 153, 287–312.

Doval, L. and V. Skreta (2018a): "Constrained Information Design: Toolkit," *arXiv preprint arXiv:1811.03588.*

——— (2018b): "Mechanism Design with Limited Commitment," *arXiv preprint arXiv:1811.03579*.

Dworczak, P. (2017): "Mechanism Design with Aftermarkets: Cutoff Mechanisms," *Available at SSRN 2859206*.

Ehrlich, D., I. Guttman, P. Schönbach, and J. Mills (1957): "Postdecision Exposure to Relevant Information," *The Journal of Abnormal and Social Psychology*, 54, 98.

Epstein, L. and Y. Halevy (2019): "Hard-to-interpret Signals," *Working paper*.

Epstein, L. G. and M. Schneider (2007): "Learning Under Ambiguity," *The Review of Economic Studies*, 74, 1275–1303.

——— (2008): "Ambiguity, Information Quality, and Asset pricing," *The Journal of Finance*, 63, 197–228.

Eyster, E., S. Li, and S. Ridout (2021): "A Theory of Ex Post Rationalization," *arXiv preprint arXiv:2107.07491*.

Festinger, L. (1957): *A Theory of Cognitive Dissonance*, Stanford University Press.

Frey, D. and M. Rosch (1984): "Information Seeking After Decisions: The Roles of Novelty of Information and Decision Reversibility," *Personality and Social Psychology Bulletin*, 10, 91–98.

Frey, D. and D. Stahlberg (1986): "Selection of Information after Receiving More or Less Reliable Self-Threatening Information," *Personality and Social Psychology Bulletin*, 12, 434–441.

Frey, D. and R. A. Wicklund (1978): "A Clarification of Selective Exposure: The Impact of Choice," *Journal of Experimental Social Psychology*, 14, 132–139.

Galperti, S. and J. Perego (2019): "Belief Meddling in Social Networks: An Information-design Approach," *Available at SSRN 3340090*.

Gensbittel, F., L. Renou, and T. Tomala (2015): "Comparisons of Ambiguous Experiments," *HEC Paris Research Paper No. ECO/SCD-2015-1074*.

Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

Golman, R., D. Hagmann, and G. Loewenstein (2017): "Information Avoidance," *Journal of Economic Literature*, 55, 96–135.

Hilbert, L. P., M. K. Noordewier, and W. W. van Dijk (2022): "The Prospective Associations Between Financial Scarcity and Financial Avoidance," *Journal of Economic Psychology*, 88, 102459.

Ichihashi, S. (2019): "Limiting Sender's Information in Bayesian Persuasion," *Games and Economic Behavior*, 117, 276–288.

Jakobsen, A. M. (2021): "An Axiomatic Model of Persuasion," *Econometrica*, 89, 2081–2116.

Jonas, E., S. Schulz-Hardt, D. Frey, and N. Thelen (2001): "Confirmation Bias in Sequential Information Search after Preliminary Decisions: An Expansion of Dissonance Theoretical Research on Selective Exposure to Information." *Journal of Personality and Social Psychology*, 80, 557.

Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Karlsson, N., G. Loewenstein, and D. Seppi (2009): "The Ostrich Effect: Selective Attention to Information," *Journal of Risk and Uncertainty*, 38, 95–115.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A Smooth Model of Decision Making under Ambiguity," *Econometrica*, 73, 1849–1892.

Kopylov, I. (2009): "Finite Additive Utility Representations for Preferences over Menus," *Journal of Economic Theory*, 144, 354–374.

Kopylov, I. and J. Noor (2018): "Commitments and Weak Resolve," *Economic Theory*, 66, 1–19.

Kőszegi, B. (2003): "Health Anxiety and Patient Behavior," *Journal of Health Economics*, 22, 1073–1084.

Krähmer, D. and R. Stone (2013): "Anticipated Regret as an Explanation of Uncertainty Aversion," *Economic Theory*, 52, 709–728.

Kreps, D. (2018): *Notes on the Theory of Choice*, Routledge.

Kreps, D. M. (1979): "A Representation Theorem for "Preference for Flexibility"," *Econometrica: Journal of the Econometric Society*, 565–577.

Kreps, D. M. and E. L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica: Journal of the Econometric Society*, 185–200.

Le Treust, M. and T. Tomala (2019): "Persuasion with Limited Communication Capacity," *Journal of Economic Theory*, 104940.

Lerman, C., C. Hughes, B. J. Trock, R. E. Myers, D. Main, A. Bonney, M. R. Abbaszadegan, A. E. Harty, B. A. Franklin, J. F. Lynch, et al. (1999): "Genetic Testing in Families with Hereditary Nonpolyposis Colon Cancer," *JAMA*, 281, 1618–1622.

Li, J. and J. Zhou (2016): "Blackwell's Informativeness Ranking with Uncertainty-averse Preferences," *Games and Economic Behavior*, 96, 18–29.

Liang, Y. (2019): "Learning from Unknown Information Sources," *Available at SSRN 3314789*.

Lipnowski, E., L. Mathevet, and D. Wei (2019): "Attention Management," *Available at SSRN 3161782*.

Loomes, G. and R. Sugden (1982): "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *The Economic Journal*, 92, 805–824.

——— (1987): "Some Implications of a More General Form of Regret Theory," *Journal of Economic Theory*, 41, 270–287.

Oster, E., I. Shoulson, and E. Dorsey (2013): "Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease," *American Economic Review*, 103, 804–30.

Persoskie, A., R. A. Ferrer, and W. M. Klein (2014): "Association of Cancer Worry and Perceived Risk with Doctor Avoidance: An Analysis of Information Avoidance in a Nationally Representative US Sample," *Journal of Behavioral Medicine*, 37, 977–987.

Sarver, T. (2008): "Anticipating Regret: Why Fewer Options May Be Better," *Econometrica*, 76, 263–305.

Savage, L. J. (1954): *The Foundations of Statistics*, New York: John Wiley & Sons.

Shishkin, D. and P. Ortoleva (2019): "Ambiguous Information and Dilation: An Experiment," *Working paper*.

Sicherman, N., G. Loewenstein, D. J. Seppi, and S. P. Utkus (2016): "Financial Attention," *The Review of Financial Studies*, 29, 863–897.

Stovall, J. E. (2018): "Temptation with Uncertain Normative Preference," *Theoretical Economics*, 13, 145–174.

Sugden, R. (1993): "An Axiomatic Foundation for Regret Theory," *Journal of Economic Theory*, 60, 159–180.

Takeoka, N. (2006): "Representation with a Unique Subjective Decision Tree," *working paper.*

Terkelsen, F. (1972): "Some Minimax Theorems," *Mathematica Scandinavica*, 31, 405–413.

Tsakas, E. and N. Tsakas (2018): "Noisy Persuasion," *Available at SSRN 2940681.*

Wald, A. (1950): *Statistical Decision Functions.*, Wiley.

Wei, D. (2018): "Persuasion Under Costly Learning," *Available at SSRN 3188302.*