Title
Models of Information Acquisition under Ambiguity

Author
Li, Jian

Publication Date
2012

Peer reviewed|Thesis/dissertation

Models of Information Acquisition under Ambiguity

by

Jian Li

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Chris Shannon, Chair
Associate Professor David S.Ahn
Professor Shachar Kariv
Associate Professor Andrew Lim

Spring 2012

Models of Information Acquisition under Ambiguity

Abstract

Models of Information Acquisition under Ambiguity

by

Jian Li

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Chris Shannon, Chair

This dissertation studies models of dynamic choices under uncertainty with endogenous information acquisition. In particular we are interested in exploring the interactions between ambiguity attitudes and the incentive to collect new information.

The first chapter explores the link between intrinsic preferences for information and ambiguity attitudes in settings with subjective uncertainty. We enrich the standard dynamic choice model in two dimensions. First, we introduce a novel choice domain that allows preferences to be indexed by the intermediate information, modeled as partitions of the underlying state space. Second, conditional on a given information partition, we allow preferences over state-contingent outcomes to depart from expected utility axioms. In particular we accommodate ambiguity sensitive preferences. We show that aversion to partial information is equivalent to a property of static preferences called Event Complementarity. We show that Event Complementarity and aversion to partial information are closely related to ambiguity attitudes. In familiar classes of ambiguity preferences, we identify conditions that characterize aversion to partial information.

The second chapter extends the basic model to allow for choices from non-singleton menus after partial information is revealed, and studies the value of information under ambiguity. We show that the value of information is not monotonic under ambiguity. Intrinsic aversion to partial information in the basic model is equivalent to a preference for perfect information in the extended model. Moreover, the value of information is not monotone in the degree of ambiguity aversion.

The third chapter studies the impact of ambiguity in a classic information acquisition model–the K-armed bandit problem. We consider a particular family of ambiguity averse preferences, the multiple-priors model [Marinacci, 2002]. A previous paper [Li, 2012] shows that major classic characterizations of optimal strategies in the K-armed bandit problems extend to incorporate ambiguity in the multiple-priors model. Here we explore new implications of ambiguity on the optimal incentive to experiment. First, increasing ambiguity in the unknown arm reduces the incentive to experiment, while increasing risk in the unknown arm typically increases the incentive to experiment. This suggests that ambiguity

can offer an explanation for the widely observed under-experimentation in novel technology and consumer products. Second, optimal experimentation in the multiple-priors bandit problem generally cannot be reduced to that in a classic bandit problem with an equivalent single prior. In particular, the lower envelope of the classic single-prior Gittins-Jones index for every prior lying in the multiple-priors set can be strictly higher than the generalized multiple-prior Gittins-Jones index. In one-dimensional parametric family, we identify monotonicity conditions under which this discrepancy disappears so an equivalent single prior exists.

To my parents, for their fullest understanding, encouragement, and support.

Contents

Contents **ii**

1 Preferences for Information and Ambiguity **1**
1.1 Introduction 1
1.2 Set-up 6
1.3 Intrinsic Preferences for Information 7
1.4 Ambiguity Preferences 15
1.5 Discussion: Second Order Belief Models 24

2 Multi-action Menus and Information Acquisition Problem **26**
2.1 Set-up 26
2.2 Value of Information under Ambiguity 28
2.3 Conclusion 30

3 An Application: Ambiguous Bandits **31**
3.1 Introduction 31
3.2 The Model 33
3.3 Comparative Statics 39
3.4 Existence of Equivalent Prior 41
3.5 Discussion 44

Bibliography **46**

A Proofs for Chapter 1 **50**

B Proofs for Chapter 2 **64**

C Proofs for Chapter 3 **67**

Acknowledgments

I am especially indebted to my advisor Chris Shannon, whose wonderful classes on general equilibrium and mathematical economics inspired me to work on advanced microtheory. Chris has always been an excellent mentor with her patience, generosity, and commitment to mathematical rigor in economics. I owe a lot to David Ahn for his constant input and guidance throughout the development of my dissertation, and his role in building the theory reading group and a closed community of theorists at Berkeley from which I benefited enormously. I want to express my deep gratitude to Matthew Rabin, whose behavioral economics class initiated my interest in modeling human decision making. I also want to thank Shachar Kariv and Bob Anderson for their invaluable encouragement, advice, and guidance throughout my graduate study.

I would also like to thank Haluk Ergin, Alessandro Pavan, Juan Sebastian Lleras, and David Dillenberger for useful suggestions and comments. I thank students in the theory group at Berkeley, especially Aniko Oery and Junjie Zhou, for discussions and supports. I also want to thank Glenda Oskar, Geehee Hong, Evgeny Yakovlev, Alvaro Ramos, Vinci Chow, Changcheng Song, and Xing Huang for being great classmates and friends. Finally, I thank my best friend Jing Shen, my aunt and uncle, and my lovely cousin Sasha for making California another home.

Chapter 1

Preferences for Information and Ambiguity

1.1 Introduction

This paper studies intrinsic preferences for information and the link between intrinsic information attitudes and ambiguity attitudes in settings with subjective uncertainty. To illustrate the problem that motivates these results, consider the situation an economics Ph.D. student on the job market faces in December. He has submitted many job applications, and is concerned about the possible job offers he might receive the following March. This future outcome depends on the candidate's quality and performance, as well as on the quality and performance of candidates from other schools and on funding and tastes of different employers. Suppose that in December, the candidate's optimal strategy is to maximize his own quality and performance in interviews and fly-outs, independent of the quality of other candidates or the demand at different employers. Starting in late December and early January, online forums like Economics Job Market Rumors post information on interview and fly-out schedules for different schools, which provides partial information about this uncertainty. This information has no instrumental value, as the candidate cannot condition his act on it. Yet candidates exhibit diverse preferences regarding this partial information. Some check very frequently for updates, while others avoid ever looking at this partial information.

Standard dynamic subjective expected utility (SEU) theory predicts that all students should be indifferent, as this information does not affect his optimal actions. We enrich the standard dynamic choice model in two dimensions. First, we introduce a novel choice domain that allows for preferences to be indexed by the intermediate information, modeled as partitions of the underlying state space. Second, conditional on a given information partition, we allow preferences over state-contingent outcomes to depart from expected utility axioms. In particular we accommodate ambiguity sensitive preferences. We show that aversion to partial information is equivalent to a property of static preferences called Event Complementarity. We then show that Event Complementarity and aversion to partial information are closely related to ambiguity attitudes. In familiar classes of ambiguity preferences, we identify conditions that characterize aversion to partial information.

To illustrate the connection between ambiguity attitudes and information preferences more explicitly, consider the classical Ellsberg Urn. The urn has 90 balls. 30 balls are red, and 60 balls are either green or yellow, with the exact proportion unknown. Bets are on the color of a ball drawn from the urn. In the static setting, a typical Ellsbergian decision maker (DM) strictly prefers betting on red to betting on green, but strictly prefer betting on the event that the ball is either green or yellow ($\{G, Y\}$), to betting on the event that the ball is either red or yellow ($\{R, Y\}$).

$$\begin{pmatrix} 1 & R \\ 0 & G \\ 0 & Y \end{pmatrix} \succ_0 \begin{pmatrix} 0 & R \\ 1 & G \\ 0 & Y \end{pmatrix} \text{ and } \begin{pmatrix} 1 & R \\ 0 & G \\ 1 & Y \end{pmatrix} \prec_0 \begin{pmatrix} 0 & R \\ 1 & G \\ 1 & Y \end{pmatrix}$$

In the classical Ellsberg paradox, the relative attractiveness of betting on red to green is reversed when yellow is also included as a winning state. One intuition for this reversal is the complementarity between G and Y: while the probabilities of single events $\{G\}$ and $\{Y\}$ are imprecise (ranging from 0 to $\frac{2}{3}$), the joint event $\{G, Y\}$ has a precise probability $\frac{2}{3}$. This complementarity is considered indicative of ambiguity (see for example, Epstein and Zhang [2001]).

Information can erase this complementarity and thus create ambiguity. To see this, suppose now there are two periods, and at the end of period 1, the DM will learn whether the drawn ball is yellow or not. This partial information can be described by the partition $\pi = \{\{R, G\}, \{Y\}\}$. The top event tree in Figure 1.1 illustrates the corresponding dynamic information structure. Suppose when expecting information π, the DM evaluates the dynamic bets by backward recursion: she first contemplates how she will rank acts at the end of stage 1, conditional on the realization of either event $\{R, G\}$ or event $\{Y\}$, and then aggregates these conditional preferences to form the ex-ante preferences expecting π. In this way, acts are evaluated separately for payoffs on events $\{R, G\}$ and payoffs on event $\{Y\}$, so the complementarity between G and Y is not taken into account. By partitioning the event $\{G, Y\}$ into the subevents $\{G\}$ and $\{Y\}$, information π breaks the complementarity between G and Y and creates ambiguity. On the other hand, if the DM is not told anything at the end of stage 1, an information structure illustrated by the bottom event tree π_0 in Figure 1, this complementarity is fully taken into account. So if a DM is ambiguity averse and values this complementarity, then she will prefer event tree π_0 to event tree π and exhibit an aversion to partial information in the interim stage.

The connection we establish between ambiguity attitudes and intrinsic preferences for partial information is important for a number of reasons. From a theoretical perspective, when ambiguity aversion implies intrinsic preferences for information, then endogenous learning and information acquisition decisions can be different from those in a standard dynamic SEU model. In particular, one criticism regarding the importance of incorporating ambiguity in the long run steady state is that in a stationary environment, ambiguity could eventually

Figure 1.1. Dynamic Ellsberg Paradox.

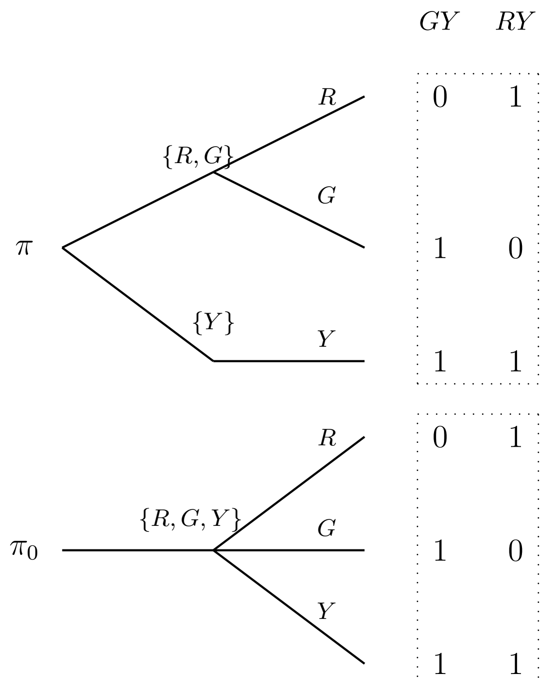

Notes: In the top event tree, the partition is $\pi = \{\{R, G\}, \{Y\}\}$. In the bottom event tree π_0, the partition is the trivial no information partition $\pi_0 = \{\{R, G, Y\}\}$.

be learnt away. If learning is endogenous and ambiguity aversion undermines the incentive to collect new information, however, then ambiguity can persist in the long-run steady state. Of more direct policy relevance, recent work illustrates the importance of ambiguity in finance and macroeconomics for providing more accurate and robust dynamic measures of risk in financial positions.[1] Our results suggest that the nature and timing of information could be an important additional component to include in the design of risk measures that account for ambiguity.

To formalize, we study a two-period model where state-dependent consequences are realized in the second period, and some partial information π, a partition of the state space S, is revealed in the first period. In particular, we relax *reduction*, that the DM is indifferent to the temporal resolution of uncertainty, so the DM has intrinsic preferences for the temporal resolution of uncertainty. Formally, we do so by considering preferences on the extended domain $\Pi \times \mathcal{F}$, the product space of information partitions and Anscombe-Aumann acts. The primitives are the ex-ante preferences $\succcurlyeq$ on $\Pi \times \mathcal{F}$, the underlying unconditional preferences

[1] For applications of ambiguity in finance and macroeconomics, see Epstein and Wang [1994], Hansen and Sargent [2001], Cao et al. [2005], and Ju and Miao [2010]. In addition, Epstein and Schneider [2010] survey applications of ambiguity preferences in finance, and Backus et al. [2005] survey applications of ambiguity preferences in macroeconomics. For work on dynamic risk measures under ambiguity, see Riedel [2004] and Acciaio et al. [2011] and references therein.

$\succcurlyeq_0$ that coincide with $\succcurlyeq$ on $\{\pi_0\} \times \mathcal{F}$ where $\pi_0 = \{S\}$ is the trivial no-information partition, as well as the family of conditional preferences $\{\succcurlyeq_E\}_{E\in\Sigma}$ on $\mathcal{F}$.

In Section 1.3, we give axioms under which $\succcurlyeq$ and $\{\succcurlyeq_E\}_{E\in\Sigma}$ have a cross-partition recursive representation. For a fixed partition π, ex-ante preferences $\succcurlyeq$ on $\{\pi\} \times \mathcal{F}$ are connected with conditional preferences $\{\succcurlyeq_E\}_{E\in\pi}$ through dynamic consistency (called π-Recursivity in text). Across partitions, we characterize an updating rule that ensures all conditional preferences $\{\succcurlyeq_E\}_{E\in\Sigma}$ are derived from the same underlying unconditional preferences $\succcurlyeq_0$. In this way, ex-ante preferences across partitions are generated by the same $\succcurlyeq_0$ and thus reflect consistent beliefs about events in S.

Under this recursive representation, we establish the equivalence between aversion to partial information in the ex-ante preferences $\succcurlyeq$ and a property on the unconditional preferences $\succcurlyeq_0$ called Event Complementarity. We show that Event Complementarity captures the intuition of complementarity in the Ellsberg example and thus the concept of ambiguity aversion. In Section 1.4, we further explore the intersection between Event Complementarity (and thus preferences for partial information) and popular models of ambiguity preferences. We find that for maxmin expected utility (MEU) [Gilboa and Schmeidler, 1989] and Choquet expected utility (CEU) [Schmeidler, 1989], there is a tight connection between ambiguity aversion (loving) and aversion (attraction) to partial information. For the more general class of variational preferences [Maccheroni et al., 2006a], this connection is more delicate. For variational preferences, we identify a condition on the cost function that characterizes aversion to partial information. We also identify joint conditions on the cost function and acts that characterize local aversion to partial information at a particular act. Finally, we show that for multiplier preferences [Hansen and Sargent, 2001, Strzalecki, 2011], ex-ante preferences exhibit partial information neutrality.

This paper makes several novel contributions. First, we identify a connection between ambiguity attitudes and preferences for partial information, which is of both theoretical and applied interest. Second, this paper introduces a model of dynamic ambiguity preferences across different information structures, and reconciles the well-known tension between dynamic consistency and ambiguity preferences through relaxing reduction.[2] Third, this paper makes an independent contribution to the study of updating rules for ambiguity sensitive preferences. In particular, we provide a behavioral characterization for a simple updating rule for variational preferences.

One limitation of this work is that the behavioral characterization for updating is only well-defined for the class of translation invariant preferences. This rules out the second order belief models [Klibanoff et al., 2005, Nau, 2006, Seo, 2009], another important family of ambiguity preferences. In Section 1.5, we discuss information preferences for second order belief models.

The rest of the paper is organized as follows. The rest of Section 1.1 discusses related literature. Section 1.2 introduces the set-up. Section 1.3 axiomatizes the cross-partition recursive representation, and shows that aversion to partial information is equivalent to

[2]This point is discussed in more detail in the related literature section.

Event Complementarity. Section 1.4 further examines the link between ambiguity aversion and aversion to partial information, by studying four popular representations of ambiguity preferences. Section 1.5 discusses the second order belief models.

1.1.1 Related Literature

This paper belongs to the literature on dynamic decision making under ambiguity. Epstein and Schneider [2003] axiomatize recursive preferences over adapted consumption processes where all conditional preferences are maxmin expected utility (MEU), and find that dynamic consistency (our π-Recursivity) implies that the prior belief set has to satisfy a "rectangularity" restriction. Later work axiomatizes recursive preferences for other static ambiguity preferences and finds similar restrictions [Maccheroni et al., 2006b, Klibanoff et al., 2008].

In fact, Siniscalchi [2011] shows that within a given filtration, dynamic consistency implies Savage's Sure-Thing Principle and Bayesian updating. Together with reduction, dynamic consistency rules out modal Ellsberg preferences and thus ambiguity.[3] To allow for ambiguity, Siniscalchi studies preferences over a richer domain of decision trees, and relaxes dynamic consistency by introducing a weaker axiom called Sophistication. Together with auxiliary axioms, he proposes a general approach where preferences can be dynamically inconsistent, and the DM addresses these inconsistencies through Strotz-type Consistent Planning.

In this paper, we start from the observation that the noted tension between dynamic consistency and ambiguity relies on reduction, that is, on the assumption that the DM is indifferent to the temporal resolution of uncertainties. However, experimental evidence suggests that reduction is often violated in environments with objective risk.[4] For example, Halevy [2007] finds evidence for non-reduction of compound lotteries and ambiguity aversion, as well as a positive association between the two. In a dynamic portfolio choice experiment, Bellemare et al. [2005] find that when a DM is committed to some ex-ante portfolio, higher frequency of information feedback leads to lower willingness to invest in risky assets. In this paper, we explore how dynamic consistency and unrestricted ambiguity preferences can be reconciled by relaxing reduction.

Thus this paper is also related to a rich literature that relaxes reduction and studies intrinsic preferences for early or late resolution of uncertainty. This was initially formalized by Kreps and Porteus [1978] by introducing a novel domain of objective temporal lotteries and subsequently extended by Epstein and Zin [1989, 1991] to study asset pricing. Grant et al. [1998, 2000] link time preferences to intrinsic preferences for information. In a purely subjective domain, Strzalecki [2010] shows that even with standard discounting most models of ambiguity aversion display some preference with regard to the timing of resolution, with the notable exception of the MEU model. Motivated by experimental evidence,[5] recent

[3] See also earlier work by Epstein and LeBreton [1993].

[4] To my best knowledge, we don't have direct evidence on violation of reduction in environments with subjective uncertainty. One potential experimental design to test reduction is the Ellsberg example illustrated in the introduction.

[5] For example, Gneezy et al. [2003], Haigh and List [2005], and Bellemare et al. [2005].

work studies preferences for one-shot versus gradual resolution of (objective) uncertainty. In the domain of objective two-stage compound lotteries,[6] Dillenberger [2010] identifies a link between preferences for one-shot resolution of uncertainty and Allais-type behaviors. In their reference-dependent utility model, Koszegi and Rabin [2009] also find preferences for getting information "clumped together rather than apart." In contrast, here we identify a link between ambiguity attitudes and intrinsic preferences for partial information over subjective uncertainty.

Finally, our work is also related to the literature on consequentialist updating rules for preferences that violate Savage's Sure-Thing Principle.[7] Pires [2002] introduces a coherence property that characterizes the prior-by-prior Bayesian updating rule for MEU preferences. Eichberger et al. [2007] then apply this coherency property to characterize full Bayesian updating for Choquet expected utility (CEU) preferences. Here we apply this property to general translation invariant preferences to connect unconditional and conditional preferences. We then show that this characterizes a simple updating rule for variational preferences, which nests previous results for Bayesian updating in the MEU and multiplier preferences cases.

1.2 Set-up

1.2.1 Preliminaries

Subjective uncertainty is modeled by a finite set S of states of the world, with elements $s \in S$, describing all contingencies that could possibly happen. Let Σ be the power set of S. $\Delta(S)$ is the set of all probabilities on S. For any $E \subseteq S$, $\Delta(E)$ denotes the set of probabilities on (S, Σ) such that $p(E) = 1$.

Z is the set of deterministic consequences. We assume that Z is a separable metric space. Let $X = \Delta(Z)$, the set of all objective lotteries over Z, endowed with the weak topology. An act $f : S \to X$ is a mapping that associates to every state a lottery in X.

Let $\mathcal{F}$ be the set of all such acts, endowed with the product topology. An act f is *constant* if there is some $x \in X$ such that $f(s) = x, \forall s$; in this case f is identified with x. For all $f, g \in \mathcal{F}, E \in \Sigma$, fEg denotes the act such that $(fEg)(s) = f(s)$ if $s \in E$, and $(fEg)(s) = g(s)$ if $s \notin E$. For any $f, g \in \mathcal{F}, \alpha \in (0, 1)$, $\alpha f + (1 - \alpha)g$ denotes the pointwise mixture of f and g: $(\alpha f + (1 - \alpha)g)(s) = \alpha f(s) + (1 - \alpha)g(s)$.

Let $B(S)$ be the space of all real-valued functions on S, endowed with the sup-norm. For any interval $K \subseteq \mathbb{R}$, $B(S, K)$ denotes the subspace of functions that take values in K.

Partial information is a partition of S. A generic partition is denoted $\pi = \{E_1, \ldots, E_n\}$, where the sets E_i are nonempty and pairwise disjoint, $E_i \in \Sigma$ for each i, and $\cup_{i=1}^n E_i = S$. Let Π be the set of all such partitions. In particular, $\pi_0 = \{S\}$ denotes the coarsest

[6]Segal [1990] was the first to study two-stage compound lotteries without reduction.

[7]Alternatively, Hanany and Klibanoff [2007, 2009] relax consequentialism, and characterize dynamically consistent updating rules for ambiguity preferences. They use a weaker notion of dynamic consistency than ours.

Figure 1.2. Information Partitions of $S = \{s_1, s_2, s_3\}$.

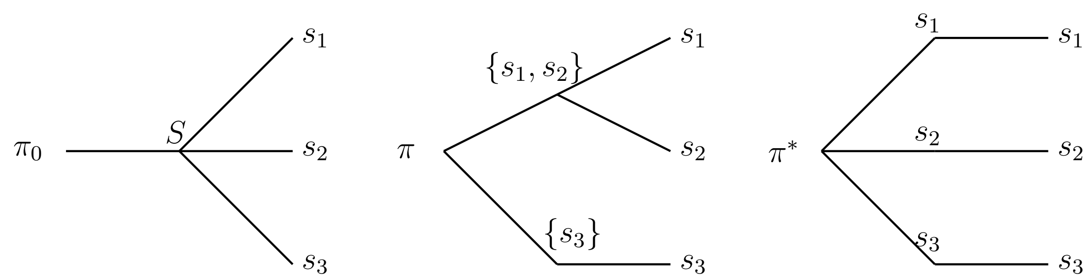

partition, capturing the case when no information is learned in the intermediate stage, and $\pi^* = \{\{s_1\}, \ldots, \{s_{|S|}\}\}$ denotes the finest partition, capturing the case when all relevant uncertainties are resolved in the intermediate stage.

Finally, for all π, let $\mathcal{F}_\pi$ be the subset of π-measurable acts in $\mathcal{F}$.

1.3 Intrinsic Preferences for Information

In this section, we show that ambiguity aversion is closely related to intrinsic information aversion. We first focus on the value of decision problems when menus are singletons, so the domain of preferences is $\Pi\times\mathcal{F}$. We develop a dynamic model of ambiguity averse preferences which retains recursivity but relaxes reduction, so information could potentially affect the evaluation of a single act. The extension to multi-action menus will be studied in the next section.

Formally, suppose the DM has ex-ante preferences $\succcurlyeq$ over $\Pi \times \mathcal{F}$.[8] Then $(\pi, f) \succcurlyeq (\pi', g)$ means that the DM prefers act f (or equivalently, the singleton menu $\{f\}$) when anticipating information π, to act g when anticipating information π'. For given information π, upon learning that the state s lies in event E in the intermediate stage, the DM updates her prior preferences $\succcurlyeq$ to E-conditional preferences $\succcurlyeq_E$. We assume that the conditional preferences $\succcurlyeq_E$ depend only on the event E but not on π, so for each E, conditional preferences $\succcurlyeq_E$ are defined on $\mathcal{F}$.[9] We also denote by $\succcurlyeq_\pi$ the restriction of $\succcurlyeq$ to $\{\pi\} \times \mathcal{F}$, interpreted as the DM's ex-ante preferences over $\mathcal{F}$ when expecting information π. Thus $\succcurlyeq$ and $\{\succcurlyeq_E\}$ are the primitive preferences of our model.

[8] We endow Π with the discrete topology, and put the product topology on $\Pi \times \mathcal{F}$.

[9] In a two period model, there is no further information to expect after some event in π is realized, so it is reasonable to have conditional preferences defined only on $\mathcal{F}$.

We look for a dynamic model of preferences over $\Pi \times \mathcal{F}$ that satisfies two criteria. First, within a given partition $\pi = \{E_1, E_2, \cdots, E_n\}$, $\succsim_\pi$ and $\{\succsim_{E_i}\}_{i=1}^n$ satisfy a recursive relation, in the following sense. For any act f, construct another act f' by replacing f on each E_i by a constant act x_i, where $x_i \sim_{E_i} f$. So $f'(s) = x_i$ if $s \in E_i$, for all i. Recursivity requires that $f \sim_\pi f'$. Second, across two different information partitions π and π', $\succsim_\pi$ and $\succsim_{\pi'}$ are related by a unifying unconditional preference relation over $\mathcal{F}$. That is, there exists an unconditional preference relation $\succsim_0$ over $\mathcal{F}$ such that all conditional preferences $\{\succsim_E\}_{E\in\Sigma}$ are updated from $\succsim_0$. Thus if we observe any difference between $\succsim_\pi$ and $\succsim_{\pi'}$, it is due to differences in π and π' rather than ex-ante beliefs.

1.3.1 Recursive Model

In this section we impose axioms on $\{\succsim_\pi\}_{\pi\in\Pi}$ and $\{\succsim_E\}_{E\in\Sigma}$ that characterize the folding back evaluation procedure.

First we impose common basic technical axioms on $\succsim_\pi$ and $\succsim_E$, for each $\pi \in \Pi$ and $E \in \Sigma$. For convenience we group them together as Axiom 1.

Axiom 1.
1. (Continuity) For all π, E, f, $\{g \in \mathcal{F} : g \succsim_\pi f\}$, $\{g \in \mathcal{F} : f \succsim_\pi g\}$, $\{g \in \mathcal{F} : g \succsim_E f\}$, and $\{g \in \mathcal{F} : f \succsim_E g\}$ are closed.
2. (Monotonicity) For all $\pi, E \in \Sigma$, if $f(s) \succsim_\pi (\succsim_E) g(s)$, $\forall s$, then $f \succsim_\pi (\succsim_E) g$.
3. (Non-degeneracy) For all π, $f \succ_\pi g$ for some $f, g \in F$. Similarly, $\forall E \in \Sigma$, $f \succ_E g$ for some $f, g \in F$.

Axiom 2 (Stable Risk Preferences)**.** For all π, E, $\succsim_\pi$ and $\succsim_E$ agree on constant acts.

Lemma 1.1. *Under Continuity and Stable Risk Preferences, $\succsim$ is a continuous preference relation on $\Pi \times \mathcal{F}$.*

Proof. See appendix. □

Within a fixed partition $\pi = \{E_1, \cdots, E_n\}$, we impose π-recursivity to link prior preferences $\succsim_\pi$ and conditional preferences $\{\succsim_{E_i}\}_{i=1}^n$. This is similar to the Dynamic Consistency axiom in Epstein and Schneider [2003] and Maccheroni et al. [2006b], simplified to two periods.

Axiom 3 (π-Recursivity)**.** For any π, $E \in \pi$, and $f, g, h \in \mathcal{F}$,

$$f \succsim_E g \Leftrightarrow fEh \succsim_\pi gEh$$

If all $\succsim_\pi$ satisfy π-Recursivity, then all conditional preferences $\{\succsim_E\}_{E\in\Sigma}$ satisfy *Consequentialism*, that is, $\forall f, g, h$, $\forall E$, $fEg \sim_E fEh$.[10] Intuitively, this says that outcomes in

[10] To see this, let $fEg = f'$ and $fEh = g'$. Then $f'Ef = g'Ef = f$. For $\pi = \{E, E^c\}$, $f'Ef \sim_\pi g'Ef$, and by π-Recursivity, $f' \sim_E g'$.

states outside E do not affect E-conditional preferences $\succcurlyeq_E$. We will return to this when discussing learning rules.

If an act f is π-measurable, then in both (π, f) and (π^*, f), all uncertainties about f are resolved in the first stage. So the additional information in π^* relative to that in π should not affect the evaluation of f. This idea is reflected in the following axiom.

Axiom 4 (Indifference to Redundant Information)**.** For all π, $f \in \mathcal{F}_\pi$, $(\pi, f) \sim (\pi^*, f)$.

The last axiom, Time Neutrality, abstracts information preferences from any effect due to preferences for early or late resolution of uncertainty, which is orthogonal to the information preferences of interest here.

Axiom 5 (Time Neutrality)**.** For all f, $(\pi^*, f) \sim (\pi_0, f)$.

Time Neutrality implies that $\succcurlyeq_{\pi^*} = \succcurlyeq_{\pi_0}$, and both can be viewed as the unconditional preferences over acts, denoted by $\succcurlyeq_0$ in the following text. In the next subsection, we specify how all conditional preferences are updated from a unifying unconditional $\succcurlyeq_0$, ensuring all $\succcurlyeq_\pi$ represent the same ex-ante belief.

For a fixed $\pi = \{E_1, \cdots, E_n\}$, we define the conditional certainty equivalent mapping $c(\cdot|\pi) : \mathcal{F} \to \mathcal{F}_\pi$, as follows:

$$c(f|\pi) = \begin{pmatrix} c(f|E_1) & E_1 \\ c(f|E_2) & E_2 \\ \cdots & \\ c(f|E_n) & E_n \end{pmatrix}$$

where for each i, $c(f|E_i) \in X$, and $c(f|E_i) \sim_{E_i} f$. That is, $c(f|E_i)$ is the certainty equivalent of f conditional on E_i. Existence is guaranteed by Continuity and Monotonicity of each $\succcurlyeq_{E_i}$, as proved in Lemma A.1 in the appendix.

Recall that $\succcurlyeq$ is the ex-ante preference over $\Pi \times \mathcal{F}$, while for every π, $\succcurlyeq_\pi$ is the restriction of $\succcurlyeq$ to $\{\pi\} \times \mathcal{F}$. For an interval $K \subseteq \mathbb{R}$, $B(S, K)$ is the space of functions on S with range K. For any $k \in K$, denote by $\bar{k}$ the corresponding constant function in $B(S)$ taking value k. For any ξ and ϕ in $B(S, K)$, and any event $E \in \Sigma$, $\xi E \phi$ denotes the function such that $(\xi E \phi)(s) = \xi(s)$ if $s \in E$, and $(\xi E \phi)(s) = \phi(s)$ if $s \notin E$. For a functional $I : B(S, K) \to \mathbb{R}$, we say I is *monotone* if $\forall \xi, \phi \in B(S, K)$, $\xi \geq \phi$ implies $I_0(\xi) \geq I_0(\phi)$, and *strongly monotone* if in addition $\xi > \phi$ implies $I_0(\xi) > I_0(\phi)$. We say I is *normalized* if $I(\bar{k}) = k$ for all $k \in K$. Finally, we say I is *translation invariant* if $I(\xi + \bar{k}) = I(\xi) + k$ for all $\xi \in B(S, K)$ and $k \in K$ such that $\xi + \bar{k} \in B(S, K)$.

Lemma 1.2. *For preferences $\succcurlyeq$ and $\{\succcurlyeq_E\}_{E \in \Sigma}$ that are continuous and monotone, the following statements are equivalent:*

1. *$\{\succcurlyeq_\pi\}_{\pi \in \Pi}$ and $\{\succcurlyeq_E\}_{E \in \Sigma}$ satisfy π-Recursivity, Independence of Redundant Information, and Time Neutrality.*

2. *There exists a continuous function* $u : X \to \mathbb{R}$*, and a continuous, monotone, and normalized function* $I_0 : B(\Sigma, u(X)) \to \mathbb{R}$ *such that for each* π*,* $\succsim_\pi$ *can be represented by* $V(\pi, \cdot) : \mathcal{F} \to \mathbb{R}$*, where*
$$V(\pi, f) = I_0(u \circ c(f|\pi))$$
and $c(\cdot|\pi) : \mathcal{F} \to \mathcal{F}_\pi$ *is the conditional certainty equivalent mapping.*

Using Axioms 1-5, preferences $\succsim_\pi$ and $\{\succsim_E\}_{E\in\pi}$ satisfy π-Recursivity, under which the value of an act f expecting information π can be computed by a folding back procedure. For each event $E_i \in \pi$, replace f on E_i by its conditional certainty equivalent. The constructed act $c(f|\pi)$ is π-measurable, thus could be evaluated by the unconditional preferences $\succsim_0$, and
$$(\pi, f) \succsim (\pi', g) \Leftrightarrow c(f|\pi) \succsim_0 c(g|\pi')$$
Therefore, the ex-ante preferences $\succsim$ are dictated by the conditional preferences $\{\succsim_E\}_{E\in\Sigma}$ and unconditional preferences $\succsim_0$.

For any π, let $B(\pi, u(X))$ denote all the π-measurable functions in $B(S, u(X))$.

1.3.2 Updating Translation Invariant Preferences

In this subsection, we characterize an updating rule that specifies how the conditional preferences $\{\succsim_E\}_{E\in\Sigma}$ are derived from unconditional preferences $\succsim_0$. In this way, for two different information partitions π and π', $\succsim_\pi$ and $\succsim_{\pi'}$ are related by the same unconditional $\succsim_0$ and thus have the same underlying beliefs about events in S. Thus any difference between $\succsim_\pi$ and $\succsim_{\pi'}$ is due to differences in information partitions π and π' rather than ex-ante beliefs. In particular, to accommodate ambiguity sensitive $\succsim_0$, we look for an updating rule that (i) requires that each $\succsim_E$ satisfies Consequentialism, so outcomes on states outside E does not affect $\succsim_E$; (ii) does not exclude a preference for hedging in $\succsim_0$.

It does not make sense to discuss conditional preferences $\succsim_E$ if event E has "probability zero". We call an event E is Savage $\succsim_0$-non-null if it is not the case that $fEh \sim_0 gEh$ for all $f, g, h \in \mathcal{F}$. For simplicity, we require that for every event E in Σ is $\succsim_0$-non-null. For the purpose of updating ambiguity preferences, we need a stronger notion of non-null events.[11] Here we ensure every event is non-null for $\succsim_0$ by imposing a strong monotonicity axiom on $\succsim_0$.

Axiom 6 (Strong Monotonicity)**.** $\forall f, g \in \mathcal{F}$, if $f(s) \succsim_0 g(s)$ for all $s \in S$, then $f \succsim_0 g$. If in addition one of the preference rankings is strict, then $f \succ_0 g$.

Bayesian updating is the universal updating rule in Savage's SEU theory. The unconditional preference is represented by an expected utility functional with respect to some subjective belief p, and the conditional preference on E is represented by an expected utility

[11] For a detailed discussion of the relationship between a Savage $\succsim_0$-non-null event and the stronger condition we need, see Appendix A.2.

functional with respect to the Bayesian posterior $p(\cdot|E)$. Behaviorally, $\succcurlyeq_E$ is derived from $\succcurlyeq_0$ by[12]

$$f \succcurlyeq_E g \Leftrightarrow fEh \succcurlyeq_0 gEh \text{ for some } h$$

We refer to this as Bayesian Updating in the rest of the paper. In Savage's theory, $\succcurlyeq_E$ is well-defined because $\succcurlyeq_0$ satisfies the Sure-Thing Principle (STP): for all f, g, h, h',

$$fEh \succcurlyeq_0 gEh \Leftrightarrow fEh' \succcurlyeq_0 gEh'$$

The Sure-Thing Principle requires that $\succcurlyeq_0$ is separable across events, which rules out a preference for hedging and Ellsberg-type preferences. This condition clearly is too strong for our purposes. Instead, we consider a weaker condition, called Conditional Certainty Equivalent Consistency. This condition requires that a constant act x is equivalent to an act f conditional on E if and only if x is also unconditionally equivalent to fEx, the act that gives f for states in E, and x for states outside E.

Axiom 7 (Conditional Certainty Equivalent Consistency)**.** $\forall f \in \mathcal{F}, x \in X$, $\forall E \in \Sigma$,

$$f \sim_E x \Leftrightarrow fEx \sim_0 x$$

Conditional Certainty Equivalent Consistency weakens Bayesisan Updating by restricting g and h to be a constant act x and considering only indifference relations. In particular, Bayesian Updating imposes two properties. First, $\succcurlyeq_0$ and $\succcurlyeq_E$ are dynamically consistent: if f and g agree outside event E, then f is preferred to g conditional on E if and only if f is preferred to g unconditionally. Second, $\succcurlyeq_E$ satisfies consequentialism: if f and g agree on event E, then f is equivalent to g conditional on E. It is straightforward to verify that under Conditional Certainty Equivalent Consistency, consequentialism is retained but not dynamic consistency.

Just as Savage's Bayesian Updating is not well-defined unless $\succcurlyeq_0$ satisfies the STP, we also need to impose some structural assumption on $\succcurlyeq_0$ to ensure that Conditional Certainty Equivalent Consistency is well-defined. The property needed is translation invariance of the corresponding aggregating functional I_0. The behavioral axiom that characterizes translation invariance is Maccheroni et al. [2006a]'s Weak Certainty Independence.[13]

Axiom 8 (Weak Certainty Independence)**.** For all $f, g \in \mathcal{F}$, $x, y \in X$, and $\alpha \in (0, 1)$,

$$\alpha f + (1-\alpha)x \succcurlyeq_0 \alpha g + (1-\alpha)x \Rightarrow \alpha f + (1-\alpha)y \succcurlyeq_0 \alpha g + (1-\alpha)y$$

Intuitively, Weak Certainty Independence of $\succcurlyeq_0$, and thus translation invariance of I_0, requires that the indifference curves in the space of utility profiles are parallel when moved

[12] See, for example, Kreps [1988, chap. 9].

[13] By Maccheroni et al. [2006a]'s Lemma 28,Weak Certainty Independence, Monotonicity, Continuity, and Non-degeneracy of $\succcurlyeq_0$ is equivalent to $\succcurlyeq_0$ can be represented by an affine risk utility u and normalized, monotone, and translation invariant functional aggregator I_0.

along the certainty line. Ambiguity preferences that satisfy translation invariance include MEU, CEU, and variational preferences. As mentioned in the discussion of related literature, Conditional Certainty Equivalence Consistency has been used by Pires [2002] to characterize prior-by-prior updating for MEU, and by Eichberger et al. [2007] to characterize a generalized Bayes rule for CEU. In Section 1.4, we characterize a simple update rule for variational preferences using this axiom.

We show that if $\succsim_0$ satisfies Weak Certainty Independence and Strong Monotonicity, then conditional preferences are well-defined. Thus only knowledge about $\succsim_0$ is needed to calculate the conditional certainty equivalent, and thus pin down the conditional preferences $\succsim_E$ for all E. Moreover, when combined with axioms characterizing recursiveness in the previous subsection, knowing $\succsim_0$ is sufficient to characterize $\succsim_\pi$ for all π. Below is a formal definition.

Definition 1.1. We say $\succsim$ on $\Pi \times \mathcal{F}$ and $\succsim_E$ on $\mathcal{F}$ *have a cross-partition recursive representation* (u, I_0) if

1. There exists a continuous, non-constant, and affine $u : X \to \mathbb{R}$, and a continuous, strongly monotone, normalized, and translation invariant $I_0 : B(S, u(X)) \to \mathbb{R}$ such that
$$f \succsim_0 g \Leftrightarrow I_0(u \circ f) \geq I_0(u \circ g)$$
2. For all $E \in \Sigma$, $\succsim_E$ is represented by $V_E : \mathcal{F} \to \mathbb{R}$, where $V_E(f)$ is the unique solution to
$$k = I_0((u \circ f)E\bar{k})$$
3. $\succsim$ is represented by $V : \Pi \times \mathcal{F} \to \mathbb{R}$, where
$$V(\pi, f) = I_0(V_0(f|\pi))$$
and
$$V_0(f|\pi) = \begin{pmatrix} V_{E_1}(f) & E_1 \\ V_{E_2}(f) & E_2 \\ \dots & \\ V_{E_n}(f) & E_n \end{pmatrix}$$

In this case, we also say $\succsim$ *is recursively generated by* $\succsim_0$.

Theorem 1.1. *The following statements are equivalent:*

1. *a)* $\{\succsim_\pi\}_{\pi \in \Pi}$ *and* $\{\succsim_E\}_{E \in \Sigma}$ *are continuous and monotone, satisfy* π*-Recursivity, Independence of Redundant Information, Time Neutrality, and Stable Risk Preferences;*

   b) $\succsim_0$ *satisfies Weak Certainty Independence and Strong Monotonicity;* $\succsim_0$ *and* $\{\succsim_E\}_{E \in \Sigma}$ *satisfy Conditional Certainty Equivalent Consistency.*

2. $\succcurlyeq$ *and* $\succcurlyeq_E$ *have a cross-partition recursive representation with* (u, I_0).

Moreover, if two affine functions u *and* u' *both represent* $\succcurlyeq_0$ *on* X*, then there exists* $a > 0, b \in \mathbb{R}$ *such that* $u' = au + b$. *For a given* u, I_0 *is unique.*

Proof. See appendix. □

1.3.3 Intrinsic Aversion to Partial Information

In this subsection, we define aversion to partial information as a property of the cross-partition preference $\succcurlyeq$. Then we show that under our recursive representation, aversion to partial information is equivalent to a property of $\succcurlyeq_0$ called Event Complementarity. We study the relationship between Event Complementarity and ambiguity aversion. In the next section, we consider familiar models of ambiguity preferences, and study the connection among ambiguity aversion, Event Complementarity, and aversion to partial information.

Definition 1.2. We say $\succcurlyeq$ exhibits *aversion to partial information at act* f if $(\pi_0, f) \succcurlyeq (\pi, f)$ for all π. We say $\succcurlyeq$ exhibits *aversion to partial information* if $\succcurlyeq$ exhibits aversion to partial information at all acts.

Attraction to partial information and information neutrality are defined analogously.

This definition of aversion to partial information is similar to Preferences for One-Shot Resolution of Uncertainty in Dillenberger [2010], and preferences to get information "clumped together rather than apart" as in Koszegi and Rabin [2009]. Our definition only requires that the DM prefers no information π_0 to any information π. This is weaker than the notion of information aversion defined in Grant et al. [1998] and Skiadas [1998], which requires that coarser information is always preferred to finer information.[14] If the DM exhibits aversion to partial information at all acts and obeys Time Neutrality, then $(\pi_0, f) \sim (\pi^*, f) \succcurlyeq (\pi, f)$ for all f.

In the modal Ellsberg preferences, there is complementarity between the events $\{G\}$ and $\{Y\}$ in eliminating ambiguity. The DM knows that the joint event $\{G, Y\}$ has a precise probability $\frac{2}{3}$, while each subevent $\{G\}$ or $\{Y\}$ has an imprecise probability ranging from 0 to $\frac{2}{3}$. By partitioning the event $\{G, Y\}$ into the subevents $\{G\}$ and $\{Y\}$, the information regarding whether the ball drawn is yellow or not breaks this complementarity and creates ambiguity. A DM averse to ambiguity might naturally be averse to this information. We formalize this idea as a condition on $\succcurlyeq_0$ below.

Axiom 9 (Event Complementarity)**.** For all E and f, if $fEx \sim_0 x$ for some x, then $f \succcurlyeq_0 xEf$.

Intuitively, Event Complementarity captures the following thought experiment. For a given act f and event E, first calibrate the value of f conditional on E by finding its

[14] In Grant et al. [1998], finer information corresponds to higher Blackwell's informativeness ranking.

conditional certainty equivalent, that is, the constant act x such that $fEx \sim_0 x(= xEx)$. Then replace f on E by x, that is, consider the act xEf, and compare this to the original act f. By construction, xEf and f are equivalent conditional on E, and they are identical, and hence trivially equivalent, conditional on E^c. A DM who satisfies the Sure-Thing Principle would view f and xEf as equivalent. Replacing f by its conditional certainty equivalent x on E, however, breaks the potential complementarity between the events E and E^c with respect to the act f. A strict preference $f \succ_0 xEf$ reveals a DM who values such complementarity.

Proposition 1.1. *Suppose $\succcurlyeq_0$ is represented by (u, I_0) where I_0 is translation invariant. Then $\succcurlyeq_0$ satisfies Event Complementarity if and only if for any act f and constant act x such that $fEx \sim_0 x$,*

$$I_0(u \circ f) \geq I_0(u \circ (fEx)) + I_0(0E(u \circ f - u \circ x)) \tag{1.1}$$

Proof. Fix f, x, E such that $fEx \sim_0 x$. By translation invariance of I_0,

$$I_0(u \circ (xEf)) = I_0(0E(u \circ f - u \circ x)) + u(x).$$

Since $fEx \sim_0 x$, $I_0(u \circ (fEx)) = u(x)$, thus

$$I_0(u \circ (xEf)) = I_0(0E(u \circ f - u \circ x)) + I_0(u \circ (fEx))$$

Thus

$$I_0(u \circ f) \geq I_0(u \circ xEf)$$

if and only if

$$I_0(u \circ f) \geq I_0(u \circ fEx) + I_0(0E(u \circ f - u \circ x))$$

So $f \succcurlyeq_0 xEf$ if and only if (1.1) holds . □

Inequality (1.1) describes Event Complementarity of $\succcurlyeq_0$ in terms of its utility representation (u, I_0). This gives us another way to understand this axiom. Given an act f and a constant act x such that $fEx \sim_0 x$, notice that the utility profile $u \circ f$ corresponding to f can be decomposed as follows:

$$u \circ f = u \circ (fEx) + 0E(u \circ f - u \circ x)$$

Since x is a constant act, $u \circ (fEx)$ varies only on E, and $0E(u \circ f - u \circ x)$ varies only on E^c by construction. Thus $u \circ f$ is decomposed into the sum of two utility profiles, one capturing the variation of $u \circ f$ on E and one capturing the variation of $u \circ f$ on E^c. Proposition 1.1 shows that Event Complementarity holds if and only if the value of utility profile $u \circ f$, $I_0(u \circ f)$, is greater than or equal to the sum of the values of these two pieces, $I_0(u \circ fEx) + I_0(0E(u \circ f - u \circ x))$. Notice that if I_0 is superadditive, then Event Complementarity holds. However, the converse is not generally true. This result will be useful in verifying that Event Complementarity holds in a number of classes of ambiguity preferences.

Finally, the following proposition shows that in our recursive model, aversion to partial information is equivalent to Event Complementarity.

Theorem 1.2. *Suppose $\succcurlyeq$ is recursively generated by $\succcurlyeq_0$. Then the following statements are equivalent:*

1. *$\succcurlyeq_0$ satisfy Event Complementarity.*
2. *$\succcurlyeq$ exhibits aversion to partial information.*

Proof. See appendix. □

1.4 Ambiguity Preferences

In this section, we investigate further the link between ambiguity aversion and aversion to partial information. In particular, we examine whether partial information aversion is implied by ambiguity aversion for four familiar classes of translation invariant ambiguity preferences: MEU, multiplier preferences, variational preferences, and CEU. Another popular class of ambiguity preferences, the second order belief model, does not satisfy translation invariance and thus is not captured by our model. We defer discussion of second order belief models to Section 1.5.

We first introduce the ambiguity aversion axiom:[15]

Axiom 10 (Ambiguity Aversion)**.** For all $f, g \in \mathcal{F}$ and $\alpha \in (0, 1)$,

$$f \sim_0 g \Rightarrow \alpha f + (1 - \alpha)g \succcurlyeq_0 f$$

As argued by Gilboa and Schmeidler [1989], Ambiguity Aversion captures a preference for state-by-state hedging. If $\succcurlyeq_0$ is represented by (u, I_0), and I_0 is continuous, monotone, normalized, and translation invariant, then $\succcurlyeq_0$ is ambiguity averse if and only if I_0 is concave.

1.4.1 Maxmin EU

MEU is the most popular model that captures ambiguity aversion. The static MEU model is axiomatized by Gilboa and Schmeidler [1989], and a recursive MEU model is axiomatized by Epstein and Schneider [2003].[16]

We say $\succcurlyeq_0$ has an MEU representation $(u, \mathcal{P})$ if it can be represented by a function $V_0 : \mathcal{F} \to \mathbb{R}$ of the form

$$V_0(f) = \min_{p \in \mathcal{P}} \int_S u(f) dp$$

where $\mathcal{P}$ is a closed and convex subset of $\Delta(S)$.

[15]In the literature, this axiom is usually called Uncertainty Aversion. Strictly speaking, it does not coincide with the definition of ambiguity aversion as in Ghirardato and Marinacci [2002] or Epstein [1999]. But for the four families of preferences we study, Axiom 10 implies ambiguity aversion.

[16]In contrast with our model, Epstein and Schneider [2003] assume reduction.

For any convex and closed prior set $\mathcal{P}$ and any partition π, we define the *π-rectangular hull of* $\mathcal{P}$ to be $rect_\pi(\mathcal{P}) = \{p = \sum_{i=1}^k p^i(\cdot|E_i)q(E_i)|\forall p^i, q \in \mathcal{P}\}$. The set $rect_\pi(\mathcal{P})$ is the largest set of probabilities that have the same marginal probabilities and conditional probabilities for events in π as elements of $\mathcal{P}$. By definition, $\mathcal{P} \subseteq rect_\pi(\mathcal{P})$ for any $\mathcal{P}$ and π. The set $\mathcal{P}$ is called *π-rectangular* if $rect_\pi(\mathcal{P}) = \mathcal{P}$. Whether $\mathcal{P}$ is π-rectangular is closely related to whether a DM with belief set $\mathcal{P}$ is strictly averse to partial information π. The next proposition summarizes the link between MEU preferences and aversion to partial information.

Proposition 1.2. *Suppose $\succsim$ is recursively generated by $\succsim_0$. Suppose $\succsim_0$ has a MEU representation $(u, \mathcal{P})$, and $\succsim_E$ has a MEU representation $(u, \mathcal{P}_E)$, for all $E \in \Sigma$. Then*

1. *$\succsim$ exhibits aversion to partial information at all acts.*
2. *For any partition π, there exists some act f such that $\succsim$ is strictly averse to π at f, i.e., $(\pi_0, f) \succ (\pi, f)$, if and only if $\mathcal{P}$ is not π-rectangular.*

Proof. See appendix. □

Remark 1. MEU has an intuitive interpretation as a malevolent Nature playing a zero-sum game against the DM [Maccheroni et al., 2006b]. In this interpretation, Nature has a constraint set $\mathcal{P}$, and chooses a probability in order to minimize the DM's expected utility. In our recursive model without reduction, the information π turns this into a sequential game. In period 0, Nature chooses a probability from $\mathcal{P}$ for events in π. In period 1, Nature chooses a (possibly different) probability from $\mathcal{P}$ over states for every event in π, conditional on that event. In this way, information π expands Nature's constraint set from $\mathcal{P}$ to $rect_\pi(\mathcal{P})$. On the other hand, the DM has committed ex-ante to a fixed act f. So introducing information π helps Nature and hurts the DM. Part (2) of Proposition 1.2 shows that if information strictly expands Nature's constraint set, that is, if $\mathcal{P} \subsetneq rect_\pi(\mathcal{P})$, then Nature can make the DM strictly worse off at some act.

Remark 2. Epstein and Schneider [2003] develop a recursive MEU model in which they maintain reduction. They show that $\succsim$ is dynamically consistent with respect to π if and only if $\mathcal{P}$ is π-rectangular. Part (2) of Proposition 1.2 shows that if we instead maintain dynamic consistency but relax reduction, then information neutrality at π is equivalent to π-rectangularity of $\mathcal{P}$.

Remark 3. When the prior set $\mathcal{P}$ is a singleton (so the DM has SEU), or when $\mathcal{P} = \Delta(S)$, the DM is intrinsically information neutral.

1.4.2 Multiplier Preferences

Introduced by Hansen and Sargent [2001] to capture concerns about model misspecification, and later axiomatized by Strzalecki [2011], multiplier preferences have found broad applications in macroeconomics.[17] We say $\succsim_0$ has a multiplier preferences representation (u, q, θ) if it can be represented by a function $V_0 : \mathcal{F} \to \mathbb{R}$ of the form

$$V_0(f) = \min_{p \in \Delta(S)} [\int u(f)dp + \theta R(p||q)]$$

where $q \in \Delta(S)$ is the reference probability, $R(p||q) = \int \ln \frac{p}{q} dp$ is the relative entropy between p and reference probability q, and θ is a scalar measuring the intensity of ambiguity aversion.

Proposition 1.3. *Suppose $\succsim_0$ has a multiplier preferences representation (u, q, θ), and $\succsim$ is recursively generated by $\succsim_0$. Then $\succsim$ exhibits intrinsic information neutrality.*

Proof. See appendix. □

1.4.3 Variational Preferences

Variational preferences are introduced and axiomatized by Maccheroni et al. [2006a,b]. We say $\succsim_0$ has a variational representation (u, c) if it can be represented by a function $V_0 : \mathcal{F} \to \mathbb{R}$ of the form

$$V_0(f) = \min_{p \in \Delta(S)} \int u(f)dp + c(p)$$

where $c : \Delta(S) \to [0, +\infty]$ is a convex, lower semicontinuous and grounded (there exists p such that $c(p) = 0$) function. The function c is interpreted as the cost of choosing a probability. The MEU model and multiplier preferences model are special cases of variational preferences.[18] Variational preferences are the most general class of ambiguity averse preferences that satisfy translation invariance.

We let $dom(c) = \{p : c(p) < +\infty\}$ denote the domain of c. If $u(X)$ is unbounded, then for a given u, c is the unique minimum convex, lower semicontinuous, and grounded cost function that represents $\succsim_0$.

Updating Variational Preferences

For any non-empty $E \in \Sigma$, we say $\succsim_E$ has a variational representation (u_E, c_E) if it can be represented by a function $V_E : \mathcal{F} \to \mathbb{R}$ of the form

$$V_E(f) = \min_{p_E \in \Delta(S)} \int_S u_E(f)dp_E + c_E(p_E)$$

where $c_E : \Delta(S) \to [0, +\infty]$ is a convex, lower-semi-continuous, and grounded *conditional cost function*.

[17] See Hansen and Sargent [2007] and references therein.

[18] Variational preferences have a MEU representation when c is 0 on a set $\mathcal{P}$ and $+\infty$ elsewhere, and a multiplier preferences representation when $c(p) = \theta R(p||q)$.

The next theorem shows that within the variational preferences family, Stable Risk Preferences and Conditional Certainty Equivalent Consistency characterize the following updating rule for conditional cost functions:

$$c_E(p_E) = \inf_{\{p \in \Delta(S): p(\cdot|E) = p_E\}} \frac{c(p)}{p(E)} \tag{1.2}$$

Taking the infimum over all probabilities with posterior p_E controls for any concern for model mis-specification outside event E, which is irrelevant to $\succsim_E$ due to consequentialism; normalization by $\frac{1}{p(E)}$ captures a maximum likelihood intuition: probabilities p assigning a higher probability on the event that occurred are more likely to be selected and determine c_E. Since we imposed Strong Monotonicity on $\succsim_0$, every event E is $\succsim_0$-non-null. In particular, $p(E) > 0$ for all $p \in dom(c)$. Then by Lemma A.4 in the Appendix, the infimum in (A.1) attains at some p.

Theorem 1.3. *Suppose $\succsim_0$ has a variational representation (u, c) and satisfies Strong Monotonicity. Suppose for any non-empty $E \in \Sigma$, $\succsim_E$ has a variational representation (u_E, c_E). Then the following are equivalent:*

1. *$\succsim_E$ and $\succsim_0$ satisfy Stable Risk Preferences and Conditional Certainty Equivalent Consistency.*

2. *$\succsim_E$ has a variational representation (u, c_E) such that*

$$f \succsim_E g \Leftrightarrow \min_{p_E \in \Delta(E)} \int_E u(f) dp_E + c_E(p_E) \geq \min_{p_E \in \Delta(E)} \int_E u(g) dp_E + c_E(p_E)$$

where

$$c_E(p_E) = \min_{\{p \in \Delta(S): p(\cdot|E) = p_E\}} \frac{c(p)}{p(E)}$$

Proof. See Appendix. □

This generalizes well-known updating rules for the two important subclasses of variational preferences: prior-by-prior updating in the MEU class, and Bayesian updating in the multiplier preferences class.

Corollary 1. *Suppose assumptions and Statement 1 in Theorem 1.3 hold.*

1. *If $\succsim_0$ also has a MEU representation $(u, \mathcal{P})$, then for any non-empty E, $\succsim_E$ has a MEU representation $(u, \mathcal{P}_E)$, where $\mathcal{P}_E$ is the set obtained from $\mathcal{P}$ by prior-by-prior updating, that is*

$$\mathcal{P}_E = \{p(\cdot|E) | p \in \mathcal{P}\}$$

2. *If $\succsim_0$ also has a multiplier preference representation (u, q, θ), then for any non-empty E, $\succsim_E$ has a multiplier preference representation (u, q_E, θ), where q_E is the Bayesian posterior of q.*

Proof. See Appendix. □

Variational Preferences and Preferences for Partial Information

In general, recursive variational preferences might not exhibit aversion to partial information at all acts. This can be explained by the following intuition. Similar to the MEU model, variational preferences also has the intuitive interpretation of a malevolent Nature playing a zero-sum game against the DM [Maccheroni et al., 2006b]. With variational preferences, Nature's constraint set is the domain of the cost function c, $dom(c)$. In addition, Nature has to pay a non-negative cost (or transfer) of $c(p)$ to the DM if it chooses a probability p in $dom(c)$. Nature seeks to minimize the DM's expected utility plus the transfer. In our recursive model without reduction, information π turns this into a sequential game, affecting both Nature's constraint set and how often Nature has to pay the DM a transfer. Similar to the MEU model, in period 0, Nature chooses a probability from $dom(c)$ for events in π. In period 1, Nature chooses a (possibly different) probability from $dom(c)$ over states for every event in π, conditional on that event. So information π expands Nature's constraint set from $dom(c)$ to $rect_\pi(dom(c))$. On the other hand, with information π, Nature also needs to pay a non-negative transfer to the DM at every node where it chooses a probability. The total transfer can be higher or lower than what Nature would have paid in the static game, depending on the cost function c. If the total transfer is higher, then this helps the DM. So the overall effect from information π is indeterminate. Below is an example in which when the transfer effect dominates and the DM strictly prefers information π at an act f.

Example 1.1 (Attraction to Partial Information in VP)**.** Suppose $S = \{s_1, s_2, s_3\}$. Let $u(x) = x$ (where $X = \mathbb{R}$). Consider the partition $\pi = \{\{s_1, s_2\}, \{s_3\}\}$. Let $E = \{s_1, s_2\}$. Let $\bar{p} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, and $\mathcal{P} = \{p \in \Delta(S) : p(s_i) \geq \delta, \forall i = 1, 2, 3\}$, for some $\delta \in (0, \frac{1}{5}]$.

Let $\alpha_{\bar{p}} = 0$. For all $p \in \mathcal{P} \backslash \bar{p}$, in the probability simplex illustrated by Figure 1.3, we connect $\bar{p}$ to p by a line segment and extend it to a point p' on the boundary of $\mathcal{P}$. Let α_p be the ratio of the length of line segment $\bar{p}p$ to the length of line segment $\bar{p}p'$. Consider the cost function

$$c(p) = \begin{cases} \alpha_p & \text{if } p \in P, \\ +\infty & \text{otherwise.} \end{cases}$$

Note that c is convex, lower semicontinuous, and grounded, so (u, c) characterizes some VP.

Consider the act $f = (0, 3K, 1K)$, where K is a large number in $\mathbb{R}_+$ and $K\delta > 10$. Without information, $V(\pi_0, f) = 4\delta K + 1$. Suppose the DM now gets partial information π. Then

$$V_E(f) = \min_{p_E \in \Delta(E)} 3K p_E(s_2) + \min_{p(\cdot|E) = p_E} \frac{c(p)}{p(E)} = \frac{1}{1-\delta}(3\delta K + 1)$$

$$V(\pi, f) = \min_p p(E) \frac{1}{1-\delta}(3\delta K + 1) + p(s_3)K + c(p) = 3\delta K + 1 + \delta K + 1 = 4\delta K + 2$$

Then $V(\pi, f) = 4\delta K + 2 > 4\delta K + 1 = V(\pi_0, f)$, so the DM has a strict preference for partial information π at f.

Figure 1.3. Probability Simplex

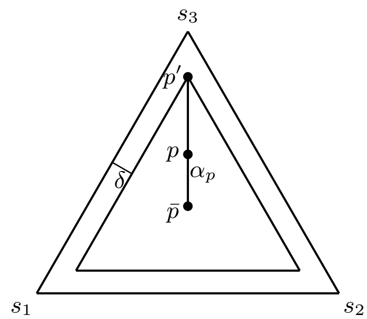

The following proposition identifies a necessary and sufficient condition on the unconditional cost function c under which aversion to partial information holds at all acts. In the zero-sum game against Nature interpretation, this condition ensures that the total transfer Nature pays under information π does not exceed that in the static game. To formalize this, we need some additional notation.

For all $p_E \in \Delta(E)$ and $p' \in \Delta(S)$, define $p_E \otimes_E p'$ by

$$(p_E \otimes_E p')(B) = p'(E)p_E(B) + p'(B \cap E^c), \forall B \in \Sigma$$

That is, in $p_E \otimes_E p'$, we substitute $p'(\cdot|E)$ by p_E for probability conditional on E, while measuring probabilities of events in E^c (including E^c) by p'.

Proposition 1.4. *Suppose $\succcurlyeq_0$ has a variational representation (u, c), and $\succcurlyeq$ is recursively generated by $\succcurlyeq_0$. Then $\succcurlyeq$ exhibits intrinsic aversion to partial information at all f if and only if for any non-empty $E \in \Sigma$,*

$$c(p) \geq \inf_{q_E \in \Delta(E)} c(q_E \otimes_E p) + p(E) \inf_{q \in \Delta(S)} \frac{c(p_E \otimes_E q)}{q(E)}, \quad \forall p, p(E) > 0$$

where p_E is the Bayesian posterior of p.

It is straightforward to verify that this condition holds for MEU and for multiplier preferences.

The above condition restricts the cost function c so that $\succcurlyeq$ exhibits partial information aversion at all acts. As shown in Example 1.1, this can be violated by some variational preferences, where attraction to partial information at some act is possible. So this condition might be too strong for some purposes.

The next proposition characterizes a joint condition on the cost function c and an act f under which $\succcurlyeq$ exhibits aversion to partial information locally at f. This does not preclude the possibility that $\succcurlyeq$ exhibits attraction to partial information at some other act g. As we will explain later, this joint condition also has an intuitive interpretation.

Proposition 1.5. *Suppose $\succsim_0$ has a variational representation (u, c), and $\succsim$ is recursively generated by $\succsim_0$. Then for any act f such that*

$$c^{-1}(0) \cap \arg\min_{p \in \Delta}[\int_S u(f)dp + c(p)] \neq \emptyset \tag{1.3}$$

$\succsim$ exhibits aversion to partial information at f.

Proof. See appendix. □

If $\succsim_0$ has MEU representation $(u, \mathcal{P})$, then the cost function is an indicator function where

$$c(p) = \delta_{\mathcal{P}}(p) = \begin{cases} 0 & \forall p \in \mathcal{P} \\ +\infty & \text{otherwise} \end{cases}$$

In this case, for any act f, $\arg\min_{p \in \Delta}[\int_S u(f)dp + c(p)] \subseteq \mathcal{P} = c^{-1}(0)$. So the result that an MEU DM is averse to partial information at all acts follows as a natural corollary of Proposition 1.5.

Condition (1.3) has an intuitive interpretation in terms of comparative ambiguity. Following the notion of comparative ambiguity aversion in Ghirardato and Marinacci [2002] and Epstein [1999], given two static preferences $\succsim_1$ and $\succsim_2$ over $\mathcal{F}$, we say $\succsim_1$ *is more ambiguity averse than* $\succsim_2$ if for all $f \in \mathcal{F}$ and $x \in X$,

$$f \succsim_1 x \Rightarrow f \succsim_2 x$$

By Maccheroni et al. [2006a] Proposition 8, if $\succsim_1$ has a variational representation (u_1, c_1) and $\succsim_2$ has a variational representation (u_2, c_2), then $\succsim_1$ is more ambiguity averse than $\succsim_2$ if and only if $u_1 \approx u_2$,[19] and $c_1 \leq c_2$ (provided $u_1 = u_2$). In the following when discussing comparative ambiguity aversion, we normalize risk utilities so that $u_1 = u_2$.[20]

We say an act f can be locally approximated by an SEU preference that is less ambiguity averse than $\succsim_0$ if there exists a preference relation $\geq'$ on $\mathcal{F}$ that admits an SEU representation

$$U'(f) = \int_S u'(f)dq$$

such that (i) $\geq'$ is less ambiguity averse than $\succsim_0$ and (ii) $V(f) = U'(f)$.

Proposition 1.6. *Suppose $\succsim_0$ has a variational representation (u, c). Condition (1.3) holds at some act f if and only if f can be locally approximated by an SEU preference that is less ambiguity averse than $\succsim_0$. In particular, if f can be locally approximated by an SEU preference that is less ambiguity averse than $\succsim_0$, then $\succsim$ exhibits aversion to partial information at f.*

[19] $u_1 \approx u_2$ if $u_1 = au_2 + b$, for some $a > 0$, $b \in \mathbb{R}$.
[20] In VP, u is unique up to positive affine transformation.

Proof. Suppose f can be locally approximated by an SEU preference $\geq'$ that is less ambiguity averse than $\succcurlyeq_0$. Let $\geq'$ be represented by U' with risk utility u' and belief $q \in \Delta(S)$. Since $\geq'$ is less ambiguity averse than $\succcurlyeq_0$, we can normalize u' so that $u = u'$. In addition, $q \in c^{-1}(0)$ by Maccheroni et al. [2006a] Lemma 32. Since $V(f) = U'(f)$,

$$V(f) = \min_{p\in\Delta}[\int_S u(f)dp + c(p)] = U'(f) = \int_S u(f)dq = \int_S u(f)dq + c(q)$$

The last equality follows from the fact that $q \in c^{-1}(0)$. So $q \in \arg\min_{p\in\Delta}[\int_S u(f)dp + c(p)]$ by definition. Together with $q \in c^{-1}(0)$, this implies that

$$c^{-1}(0) \cap \arg\min_{p\in\Delta}[\int_S u(f)dp + c(p)] \neq \emptyset$$

Thus condition (1.3) holds at f.

Now suppose there exists some $p^* \in c^{-1}(0) \cap \arg\min_{p\in\Delta}[\int_S u(f)dp + c(p)]$. Define U' by $U'(f) = \int_S u(f)dp^*$. Then by definition U' represents an SEU preference $\geq'$ that is less ambiguity averse than $\succcurlyeq_0$. Also

$$V(f) = \int_S u(f)dp^* + c(p^*) = \int_S u(f)dp^* = U'(f)$$

So f can be locally approximated by an SEU preference that is less ambiguity averse than $\succcurlyeq_0$. □

Proposition 1.7. *Suppose $\succcurlyeq_0^1$ has a variational representation (u^1, c^1) and f can be locally approximated by some SEU preference $\geq'$ that is less ambiguity averse than $\succcurlyeq_0^1$. Suppose $\succcurlyeq_0^2$ also has a variational representation (u^2, c^2), and let $\succcurlyeq^2$ be recursively generated by $\succcurlyeq_0^2$. If $\succcurlyeq_0^2$ is less ambiguity averse than $\succcurlyeq_0^1$ and more ambiguity averse than $\geq'$, then $\succcurlyeq^2$ exhibits partial information aversion at f.*

Proof. By Proposition 1.6, f can be locally approximated by an SEU preference $\geq'$ that is less ambiguity averse than $\succcurlyeq_0^1$ if and only if condition (1.3) holds. Then there exists $p^* \in c_1^{-1}(0) \cap argmin_{p\in\Delta}[\int_S u_1(f)dp + c_1(p)]$ such that $V_1(f) = \int_S u_1(f)dp^* + c_1(p^*)$, and $c_1(p^*) = 0$. By definition, $\succcurlyeq_0^2$ is less ambiguity averse than $\succcurlyeq_0^1$ if and only if $u_1 = u_2$ and $c_2 \geq c_1$. Since $\succcurlyeq_0^2$ is more ambiguity averse than $\geq'$, $u_2 = u'$ and $p^* \in c_2^{-1}(0)$. Let $u = u_1 = u_2 = u'$. Therefore:

$$\int_S u(f)dp^* + c_2(p^*) = \int_S u(f)dp^* + c_1(p^*) \leq \int_S u(f)dp + c_1(p) \leq \int_S u(f)dp + c_2(p), \forall p \in \Delta(S)$$

The first inequality follows from the fact that $p^* \in \arg\min_{p\in\Delta}[\int_S u_1(f)dp + c_1(p)]$, and the second from $c_1 \leq c_2$. Thus $p^* \in \arg\min_{p\in\Delta}[\int_S u(f)dp + c_2(p)]$. So

$$\arg\min_{p\in\Delta}[\int_S u(f)dp + c_2(p)] \cap c_2^{-1}(0) \neq \emptyset$$

and by Proposition 1.5, $\succcurlyeq^2$ exhibits aversion to partial information at f. □

1.4.4 Choquet EU

Finally, we look at the CEU model axiomatized by Schmeidler [1989]. The CEU model is of particular interest because it allows for both ambiguity averse and ambiguity loving preferences, so this provides a framework for studying the relationship between information preferences and ambiguity attitudes more generally.

We say $\succcurlyeq_0$ has a CEU representation (u, ν) if it can be represented by a function $V_0 : \mathcal{F} \to \mathbb{R}$ of the form

$$V_0(f) = \int u(f) d\nu$$

where $\nu : \Sigma \to [0, 1]$ is a capacity, that is, $\nu(S) = 1$, $\nu(\emptyset) = 0$, and $\nu(E) \leq \nu(F)$ for all $E \subseteq F$.

If $\succcurlyeq_0$ satisfies Ambiguity Aversion, then ν is a convex capacity.[21] In this case, CEU preferences become a special case of MEU preferences, with the set of priors $\mathcal{P}$ being the core of the convex capacity ν.[22] So for CEU preferences, ambiguity aversion implies aversion to partial information.

For CEU preferences, we can say a bit more about the connection between ambiguity attitudes and information preferences. We can also define ambiguity loving.[23]

Axiom 11 (Ambiguity Loving)**.** For all $f, g \in \mathcal{F}$ and $\alpha \in [0, 1]$,

$$f \sim_0 g \Rightarrow f \succcurlyeq_0 \alpha f + (1 - \alpha) g$$

We show that within the CEU model, ambiguity aversion implies partial information aversion, and ambiguity loving implies partial information loving.

Proposition 1.8. *Suppose $\succcurlyeq_0$, $\{\succcurlyeq_E\}_{E \in \Sigma}$ have CEU representations, and $\succcurlyeq$ is recursively generated by $\succcurlyeq_0$.*

1. *If $\succcurlyeq_0$ satisfies Ambiguity Aversion, then $\succcurlyeq$ exhibits partial information aversion at all acts.*

2. *If $\succcurlyeq_0$ satisfies Ambiguity Loving, then $\succcurlyeq$ exhibits attraction to partial information at all acts.*

Proof. See appendix. □

[21] A capacity ν is convex if $\nu(E \cup F) + \nu(E \cap F) \geq \nu(E) + \nu(F)$ holds for all $E, F \in \Sigma$.
[22] For a convex capacity ν, its core is $\{p \in \Delta(S) | p(E) \geq \nu(E) \text{ for all } E \in \Sigma\}$.
[23] This is called "uncertainty appeal" in Schmeidler [1989].

1.5 Discussion: Second Order Belief Models

Another important class of ambiguity preferences is the second order belief model [Klibanoff et al., 2005, Seo, 2009, Nau, 2006]. We say $\succsim_0$ has a second order belief representation if

$$V(\pi_0, f) = \int_{\Delta(S)} \phi[\int_S u(f) dp_\theta] d\mu$$

where $\mu \in \Delta(\Delta(S))$ is a second order belief over the space of distributions $\Delta(S)$, and ϕ is a non-decreasing function capturing ambiguity attitude. When ϕ is smooth and concave (convex), the DM is ambiguity averse (loving).

For the second order belief models, translation invariance fails, and thus Conditional Certainty Equivalent Consistency cannot provide a well-defined update rule. Instead we adopt Bayes rule for the second order belief μ as our update rule.

Assumption 1. Suppose $\succsim_0$ has a second order belief representation $(u, \phi; \Theta, \mu)$. Then for any non-null event E, $\succsim_E$ has a second order belief representation $(u_E, \phi_E; \Theta_E, \mu_E)$ satisfying

1. Risk and ambiguity attitudes are not updated: $u_E = u$, $\phi_E = \phi$.
2. Prior by prior updating of first order belief: $\Theta_E = \{p_\theta(\cdot|E) | p_\theta \in \Theta\}$.
3. Bayes rule for second order belief:

$$\mu_E(\theta) = \frac{\mu(\theta) p_\theta(E)}{\int_\Theta p_{\theta'}(E) d\mu(\theta')} \tag{1.4}$$

In general, second order belief models exhibit no systematic relation between ambiguity aversion and information aversion, as the following example illustrates.

Example 1.2. Consider the standard three color Ellsberg urn. Let $S = \{R, G, Y\}$ and $\Theta = \{(\frac{1}{3}, \frac{2}{3}\theta, \frac{2}{3}(1-\theta)) | \theta = \frac{1}{3}, \frac{2}{3}\}$. Suppose the second order prior μ puts equal probability on $p_{\frac{1}{3}} = (\frac{1}{3}, \frac{2}{9}, \frac{4}{9})$, and $p_{\frac{2}{3}} = (\frac{1}{3}, \frac{4}{9}, \frac{2}{9})$. Assume the DM is risk neutral with $u(x) = x$, and ambiguity averse with $\phi(y) = log(y)$. Information is given by the partition $\pi = \{\{R, G\}, \{Y\}\}$. Let $E = \{R, G\}$. Suppose the above update rule captures conditional preferences, so $\mu_E(p_{\frac{1}{3}}) = \frac{5}{12}$, and $\mu_E(p_{\frac{2}{3}}) = \frac{7}{12}$. By computation we can show that the DM is strictly averse to π $(V(\pi, f) < V(\pi_0, f))$ at acts $f = (1, 0, 0)$ and $(0, 1, 1)$, and strictly loves π $(V(\pi, f) > V(\pi_0, f))$ at acts $f = (0, 1, 0)$ and $(1, 0, 1)$.

Observe that the partition $\pi' = \{\{R\}, \{G, Y\}\}$ contains only events with known probabilities. The two acts $(1, 0, 0)$ and $(0, 1, 1)$, at which the DM is strictly averse to partial information π, are measurable with respect to π' and thus unambiguous. This suggests that a DM with second order belief preferences will be averse to partial information at acts where she has local ambiguity neutrality. The next proposition formalizes this idea.

Following Definition 4 in Klibanoff et al. [2005], we say $\succsim_0$ displays (local) smooth ambiguity neutrality at act f if $V(\pi_0, f) = \phi[\int_{\Delta(S)} \int_S u(f) dp_\theta d\mu]$. In second order belief models, ambiguity aversion only implies partial information aversion at the subclass of locally ambiguity neutral acts.

Proposition 1.9. *Suppose $\succsim_0$ and $\{\succsim_E\}_{E\in\Sigma}$ are second order belief preferences, with update rule satisfying Assumption 1. If $\succsim_0$ is ambiguity averse (loving), then $\succsim$ exhibits partial information aversion (loving) at all acts where $\succsim_0$ displays (local) smooth ambiguity neutrality.*

Proof. See appendix. □

Chapter 2

Multi-action Menus and Information Acquisition Problem

2.1 Set-up

In this chapter we study decision problems with general multi-action menus. We consider a two stage information acquisition problem. The DM is endowed with some compact menu $F \subseteq \mathcal{F}$. At stage 1, the DM acquires some partial information π by paying a cost $C(\pi)$, where $C : \Pi \to \mathbb{R}$. At stage 2, she learns which event in π realizes, and chooses an action from the menu F contingent on that event. Finally, the state s realizes and the DM receives the consequence of her chosen action.

For any menu F, the information acquisition decision reflects the standard tradeoff between the cost and benefit of getting information π. The DM will choose $\pi \in \Pi$ to solve

$$\max_{\pi \in \Pi} V(\pi, F) - C(\pi)$$

where $V(\pi, F)$ is the value of the decision problem (π, F). Because the cost $C(\pi)$ is deterministic, we focus on how the value function is affected by ambiguity attitudes. For a given menu F, the DM trades off the marginal cost and benefit of getting finer information to determine the optimal partition.

For any $\pi = \{E_1, \ldots, E_n\}$, $\forall E_i$, let f_i^* be the optimal act conditional on event E_i. Ex-ante, if information π is chosen, the DM can expect to get state contingent consequence of $f^* = f_1^* E_1 f_2^* E_2 \cdots f_{n-1}^* E_{n-1} f_n^*$, and the value of decision problem (π, F) is given by $V(\pi, F) = V(\pi, f^*)$. So the information acquisition problem can be reduced to the study of $V : \Pi \times \mathcal{F} \to \mathbb{R}$, the evaluation of singleton menus, expecting intermediate information π.

Let $\mathcal{M}$ be the collection of compact subsets of $\mathcal{F}$. We want to extend preferences over information and singleton menus, $\succcurlyeq$ on $\Pi \times \mathcal{F}$, to preferences over information and menus $\succcurlyeq^+$ over $\Pi \times \mathcal{M}$. This extension is straightforward since $\succcurlyeq$ is π-recursive for each π.

To that end, for every $F \in \mathcal{M}$ and $\pi = \{E_1, \cdots, E_n\}$, define

$$F^\pi = \{f_1 E_1 f_2 E_2 \cdots E_{n-1} f_n : f_i \in F, \forall i = 1, \cdots, n\}.$$

Note $F \subseteq F^\pi \subseteq \mathcal{F}$, and $F = F^\pi$ whenever F is a singleton.

Next, for a menu F and partition π, we define its conditional certainty equivalent as

$$c(F|\pi) = \begin{pmatrix} c(F|E_1) & E_1 \\ c(F|E_2) & E_2 \\ \cdots & \\ c(F|E_n) & E_n \end{pmatrix}$$

where $c(F|E_i) \in X$ and

$$u(c(F|E_i)) = \max_{f \in F} V_0(f|E_i)$$

We define the preferences $\succsim^+$ on $\Pi \times \mathcal{M}$ as follows:

$$(\pi, F) \succsim^+ (\pi', G) \text{ if and only if } \forall g \in G^{\pi'}, \exists f \in F^\pi, (\pi, f) \succsim (\pi', g)$$

In this case we say $\succsim^+$ *is extended from* $\succsim$.

Lemma 2.1. *Suppose $V : \Pi \times \mathcal{F} \to \mathbb{R}$ represents $\succsim$. If $\succsim^+$ is extended from $\succsim$, then $\succsim^+$ is represented by $\tilde{V} : \Pi \times \mathcal{M} \to \mathbb{R}$ where*

$$\tilde{V}(\pi, F) = \max_{f \in F^\pi} V(\pi, f) = V_0(c(F|\pi))$$

Since $\tilde{V}$ and V agree on $\Pi \times \mathcal{F}$, we abuse notation a bit by using V to denote the extended function $\tilde{V} : \Pi \times \mathcal{M} \to \mathbb{R}$. Here $V(\pi, F)$ is interpreted as the value of the decision problem (π, F).

The following proposition states some comparative statics of $V(\pi, F)$.

Proposition 2.1. *1. If $F \subseteq F'$, then $V(\pi, F) \leq V(\pi, F')$.*

2. Suppose $\succsim^1$ and $\succsim^2$ are recursively generated by variational preferences $\succsim_0^1$ and $\succsim_0^2$. If $\succsim_0^1$ is more ambiguity averse than $\succsim_0^2$, then $\forall \pi, F$, $V^1(\pi, F) \leq V^2(\pi, F)$.

The proof is straightforward and thus omitted.

Part (1) of Proposition 2.1 says that the DM always weakly prefers bigger menus. This distinguishes our model from that in Siniscalchi [2011]. In Siniscalchi [2011], the DM might prefer a smaller menu due to dynamic inconsistency and desire for commitment. This suggests one way to test the two models.

Part (2) of Proposition 2.1 says that the more ambiguity averse the DM is, the less she values any information and menu pair (π, F). However, this does not say that the value of information is decreasing in the degree of ambiguity aversion. Example 2.2 below illustrates this point.

Furthermore, $V(\pi, F)$ is not monotone in information π, so more information can be strictly worse. Formally, π_2 is (strictly) more informative than π_1, denoted $\pi_2 \geq (>) \pi_1$, if the partition π_2 is (strictly) finer than the partition π_1. If $\succsim_0$ displays non-trivial ambiguity aversion, then we can find a menu F and partitions $\pi_2 > \pi_1$ such that $V(\pi_2, F) < V(\pi_1, F)$. Below is an example.

Example 2.1. Suppose $S = \{s_1, s_2, s_3\}$, and $\succsim_0$ has a MEU representation $(u, \mathcal{P})$ where $\mathcal{P} = \{p \in \Delta^3 | p(s_1) = \frac{1}{3}, p(s_3) \in [\frac{1}{6}, \frac{1}{2}]\}$. For simplicity assume risk neutrality, so $u(x) = x$. Suppose the DM faces menu $F = \{(0, 1, 1), (0.49, 0.49, 0.49)\}$. Then $V(\pi_0, F) = \frac{2}{3}$. Let $\pi = \{\{s_1, s_2\}, \{s_3\}\} > \pi_0$. The informed DM will choose $(0.49, 0.49, 0.49)$ given $\{s_1, s_2\}$, and $(0, 1, 1)$ given $\{s_3\}$. Therefore $V(\pi, F) = 0.575 < \frac{2}{3} = V(\pi_0, F)$.[1] Information hurts.

This non-monotonicity is driven by intrinsic aversion to partial. Dillenberger [2010] shows that a preference for one-shot resolution of uncertainty in two-stage compound lotteries is equivalent to a preference for perfect information in an extended model with intermediate choices. We show a similar result is also true in our model.

We say that $\succsim^+$ exhibits *a preference for perfect information* if $\forall F \in \mathcal{M}$ and $\pi \in \Pi$, $(\pi^*, F) \succsim^+ (\pi, F)$.

Proposition 2.2. *Suppose $\succsim$ is recursively generated by $\succsim_0$, and $\succsim^+$ is extended from $\succsim$. Then the following statements are equivalent:*

1. *$\succsim^+$ exhibits a preference for perfect information.*

2. *$\succsim$ exhibits partial information aversion at all acts $f \in F$.*

3. *$\succsim_0$ satisfies Event Complementarity.*

Proof. See appendix. □

2.2 Value of Information under Ambiguity

In the rest of this section, we focus on the value of acquiring information π: $\Delta V(\pi, F) := V(\pi, F) - V(\pi_0, F)$. In appendix B, we analyze the marginal value of information, $V(\pi_2, F) - V(\pi_1, F)$ for any $\pi_2 \geq \pi_1$.

Acquiring information π affects the decision problem in two ways. First, information provides a way for the DM to fine-tune her strategy: expecting to get π, she conditions her choice of optimal action on the event realized in π, so her effective menu expands from F to F^π. This captures the instrumental value of information, and is always non-negative. Second, information directly affects the DM's utility from acts, thus also has intrinsic value. The value of information π in decision problem F admits the following decomposition:

$$\begin{aligned}\Delta V(\pi, F) &= V(\pi, F) - V(\pi_0, F) \\ &= [\max_{f \in F^\pi} V(\pi, F) - \max_{f \in F} V(\pi, f)] + [\max_{f \in F} V(\pi, f) - \max_{f \in F} V(\pi_0, f)]\end{aligned}$$

The first bracketed term captures the non-negative instrumental value of information. The second bracketed term captures the intrinsic value of information. It is zero if the DM is intrinsically neutral to information, so $V(\pi, f) = V(\pi_0, f)$ for all f, and non-positive if the

[1] $\mathcal{P}(s_2 | \{s_1, s_2\}) = [\frac{1}{3}, \frac{3}{5}]$, so $(0.49, 0.49, 0.49) \succ_{\{s_1, s_2\}} (0, 1, 1)$.

DM is averse to partial information. So a DM's willingness to pay for information π is the resulting trade-off of these two components.

Next we look for conditions under which the value of information is non-negative, that is, the DM is still willing to acquire information π when it is free, regardless of ambiguity.

Let $F_0 = \arg\max_{f \in F} V(\pi_0, f)$ be the set of uninformed optimal acts.

Proposition 2.3. *For any menu F, if there exists an uninformed optimal act f_0 that is π-measurable, then $\Delta V(\pi, F) \geq 0$.*

Corollary 2. *Suppose $\succcurlyeq_0^1$ has a variational representation (u_1, c_1), and $x \in X$ is an uninformed optimal act from menu F for DM 1. If DM 2 has a variational representation (u_2, c_2) and is more ambiguity averse than DM 1, then $\Delta V^2(\pi, F) \geq 0$.*

Proposition 2.1 says that for variational preferences $\succcurlyeq_0$, $V(\pi, F)$ is decreasing in the degree of ambiguity aversion in $\succcurlyeq_0$ for all (π, F). Is the same comparative statics true for the value of information $\Delta V(\pi, F)$? The answer is no. The value of information is non-monotone in the degree of ambiguity aversion. Below is an example.

Example 2.2. Suppose DM 1 has SEU preferences with belief $p \in \Delta(S)$. DM 2 has MEU preferences with non-singleton prior set $\mathcal{P} \subsetneq \Delta(S)$, and $\mathcal{P}$ is not rectangular with respect to some partition π (therefore $\pi > \pi_0$). DM 3 has MEU preferences with prior set $\mathcal{Q} = rect_\pi(\mathcal{P})$. Assume further that these three DMs have the same risk preferences, so DM 3 is more ambiguity averse than DM 2, and DM 2 is more ambiguity averse than DM 1.

Since $\mathcal{P} \subsetneq \mathcal{Q}$, there exists $f \in \mathcal{F}$ such that $V^2(\pi_0, f) > V^3(\pi_0, f)$. Also $V^2(\pi, f) = V^3(\pi_0, f) = V^3(\pi, f)$.[2] Therefore

$$V^3(\pi, f) - V^3(\pi_0, f) > V^2(\pi, f) - V^2(\pi_0, f).$$

Increasing ambiguity aversion *increases* the value of information π in this case.

Alternatively, DM 1 is intrinsically neutral to information, so $V^1(\pi, f) = V^1(\pi_0, f)$. Therefore

$$V^1(\pi, f) - V^1(\pi_0, f) = 0 > V^2(\pi, f) - V^2(\pi_0, f).$$

Increasing ambiguity aversion *decreases* the value of information π in this case.

Finally, we end this section with an application to portfolio choice problems.

Example 2.3 (Portfolio Choice)**.** Consider the portfolio choice example in Dow and Werlang (1992). Suppose there is a risk-neutral DM with wealth W. There is a risky asset with unit price P and present value that is either high, H, or low, L. The DM has MEU preferences and believes the probability of H belongs to the interval $[\underline{p}, \bar{p}]$. For simplicity, we assume

[2] The argument is similar to that in the proof of Proposition 1.2.

the DM could choose to buy a unit of the risky asset (B), short-sell a unit of the risky asset (S), or not do anything (N). So $F = \{B, S, N\}$. The DM's optimal portfolio choice is

$$f_0^*(P) = \begin{cases} B & \text{if } \underline{p}H + (1-\underline{p})L > P; \\ N & \text{if } \bar{p}H + (1-\bar{p})L \geq P \geq \underline{\pi}H + (1-\underline{p})L; \\ S & \text{if } P > \bar{p}H + (1-\bar{p})L. \end{cases}$$

We now add an information acquisition stage before the portfolio choice. The DM can acquire a binary signal, $\pi = \{h, l\}$, which is correlated with the state of the risky asset, with $p(h|H) = p(l|L) = q > \frac{1}{2}$. We want to know if the DM will collect information π if it is costless.

Suppose the DM's uninformed optimal choice is B. Then $V(\pi_0, B) = \underline{p}H + (1-\underline{p})L - P$, and $V(\pi, B) = [\underline{p}qH + (1-\underline{p})(1-q)L + \underline{p}(1-q)H + (1-\underline{p})qL - P] = V(\pi_0, \bar{B})$. By Lemma B.1 in the appendix, π is valuable. The other two cases could be calculated similarly. Without the need to compute the informed optimal strategies and $V(\pi, F)$, we can conclude that in this portfolio choice problem the DM will want to collect information π if it is costless.

2.3 Conclusion

In this chapter, we extend the basic model in Chapter 1 to allow for choices from menus after partial information is revealed, and study the properties of the value of information under ambiguity. We show that the value of information is not monotonic under ambiguity. Intrinsic aversion to partial information in the basic model is equivalent to a preference for perfect information in the extended model.[3] Moreover, the value of information is not monotone in the degree of ambiguity aversion.

[3] This is similar to Proposition 2 in Dillenberger [2010].

Chapter 3

An Application: Ambiguous Bandits

3.1 Introduction

The K-armed bandit problem has found interesting economics applications, for it captures a trade-off between exploration (experiment to find out the arm that pays the most) and exploitation (choose the arm that pays the most according to current knowledge). It has been applied to model economic problems like job search, consumer behavior and market pricing, research and development [Bergemann and Valimaki, 2008], adoption of new technology [Bryan, 2010], and collective experimentation [Strulovici, 2010].

The classic model assumes that the decision maker has a unique prior belief about the payoff distributions. Recent applications have found that Knightian uncertainty or ambiguity, where the DM has little information and does not have a unique prior about the payoff distributions, can also be a relevant factor for these applications. For example, workers searching for a new job might not know the exact distribution of matching [Nishimura and Ozaki, 2004], farmers considering for a new technology might not know the distribution of its productivity [Bryan, 2010], investors might not know the return distribution of financial assets [Epstein and Wang, 1994].

Motivated by the concern, this paper incorporates ambiguity, by allowing the decision maker (DM) to have multiple Bayesian priors about the likelihood distribution of each arm [Marinacci, 2002], into the classic K-armed bandit problem. In every period, exactly one arm is chosen and observation generated for this arm. The DM updates her beliefs about this arm prior-by-prior based on this observation while maintains beliefs about other arms unchanged (arms are independent). We assume that the DM evaluates the random payoff stream resulting from a strategy backward recursively, applying maxmin EU criterion [Gilboa and Schmeidler, 1989] for the sum of instantaneous utility and discounted next period continuity utility period-by-period at all histories. This ensures that dynamic programming techniques can be applied.

An earlier paper, Li [2012], studies the same set-up and shows that classic characterization of the K-armed bandit problem extends under multiple-priors (MP) utility. In particular, in one-armed bandit problem where there is only an unknown arm and a safe arm, the optimal strategy is also a switching time strategy characterized by a generalized MP Gittins index. In the general case with K unknown arms, the seminal Gittins-Jones index theorem [Gittins and Jones, 1979] extends to the MP case. This highlights that it is the independence of arms rather than expected utility assumption that is driving the Gittins-Jones index theorem.

Building on findings in Li [2012], in this chapter we explore implications of ambiguity on the optimal incentive to experiment that are different from those of risk. First, comparative statics on ambiguity differs from that on risk. For a given DM (thus fixed risk and ambiguity attitudes), while increasing risk increases the incentive to experiment, increasing ambiguity reduces the incentive to experiment. This provides qualitatively different behavioral implications of risk and ambiguity that are testable in data. This also suggests that ambiguity can provide an explanation for the widely observed under-experimentation in new technology and consumer products.[1] Second, we characterize an upper bound for the multiple-prior Gittins-Jones index, as the lower envelope of the classic single-prior Gittins-Jones index for every prior lying in the multiple-priors set. We show by an example that this upper bound can be strict, and identify conditions under which this upper bound is exact. Finally, the bandit model provides an easy justification for why ambiguity can persist in the long run steady state: if information acquisition decision is endogenous and the DM has a safe arm as outside option, then the DM might stop learning at some finite time and remain ambiguous about the unknown arm.

3.1.1 Related Literature

We will not attempt to review the vast literature on the classic K-armed bandit problem.[2] As mentioned above, this paper clearly builds on a previous paper [Li, 2012]. Below we will only comment on some recent studies that incorporates ambiguity averse preferences into one-armed bandit/optimal stopping problems.

Anderson [2012] studies the one-armed bandit problem with Bernoulli distributed payoffs, and the DM has second-order belief ambiguity averse preferences by Kahn and Sarin [1988] (which captures a failure to reduce objective compound lotteries). He show theoretically and test experimentally two behavioral predictions. First, ambiguity averse agents have a lower Gittins index than ambiguity neutral agents, appearing to undervalue information for experimentation. Second, the ambiguity averse agent is also willing to pay more than ambiguity neutral agents to learn the true mean of the payoff distribution. Different from Anderson [2012], we use a multiple (Bayesian) priors model [Marinacci, 2002] to capture ambiguity aversion, and consider arbitrary bounded payoff distributions. Our comparative statics result in Section 3.3 can be viewed as a generalization of Anderson's first claim in the

[1] For empirical evidence on underexperimentation, see Anderson [2012] and references therein.

[2] Gittins et al. [2011] is an up-to-date textbook on the topic. Bergemann and Valimaki [2008] survey the economic applications of the multi-armed bandit problem.

multiple priors case. His second claim is not true in our model.[3]

Riedel [2009] studies optimal stopping problem with multiple priors. He imposes a time consistency assumption on the set of priors over the full state space and considers only a lump-sum payoff at the stopping time. He shows that the Snell envelope and optimal stopping strategy extends to the multiple-priors case in a straightforward way. In his model, a duality theorem that the multiple-prior Snell envelope equals to lower envelop of the individual single-prior Snell envelopes and a minimax identity between the utility maximizing stopping time and the minimizing prior holds. He also characterizes a multiple-priors version of Doob Decomposition of supermartingales and optional sampling theorem. Miao and Wang [2011] apply the multiple-priors model to study option exercise and investment. Different from Riedel [2009] but similar to our model, they consider recursive utilities by specifying only sets of one-step-ahead probabilities.

The main difference between our approach and that in Riedel [2009] and Miao and Wang [2011] is that we explicitly consider a multiple priors Bayesian learning process.[4] Somewhat surprisingly, with the particular Bayesian learning structure imposed, the minimax type result in Riedel [2009] fails: the multiple-prior Gittins index is not always equal to the lower envelope of the single-prior Gittins indices. This suggests a modeling conflict between time-consistency (recursive models) and conditional i.i.d. assumption.

The rest of the paper is organized as follows. Section 3.2 introduces the K-armed bandit model with multiple-priors, a recursive construction of the multiple-prior utility, and previous results [Li, 2012] on the existence of optimal strategies and characterizations in the special case of one-armed bandit problem that are necessary for later explorations. Section 3.3 explores the comparative statics of the optimal experimentation on ambiguity. Section 3.4 questions the existence of an equivalent single prior for every multiple-prior bandit problem. Section 3.5 discusses an alternative specification of the multiple-prior utility. Some of the proofs are relegated to Appendix C.

3.2 The Model

In this section, we will first introduce the model set-up and a recursive construction of the multiple-prior utility. Then we will summarize results in Li [2012] on the existence of optimal strategies, and characterization of the optimal strategies and value function by a generalized MP Gittins index in one-armed bandit problems. In Section 3.3 and 3.4, we will explore further properties of the MP Gittins index.

[3] One crucial assumption for Anderson [2012]'s Theorem 2 is that an ambiguity averse agent computes the ex-ante value from learning the true distribution (θ) by expected utility. In our paper, the agent is always ambiguity averse and adopts a multiple priors model to compute one-step-ahead utilities.

[4] Therefore every prior considered describes a conditional i.i.d. process.

3.2.1 Set-up

The set-up is similar to that in Li [2012]. We consider a classic K-armed bandit problem.[5] Time is discrete and varies over $\{1,\ldots,T\}$. Horizon T can be either finite or infinite. There are K independent arms, which can be interpreted as competing R&D projects, technologies, consumer products, jobs, etc. For each period t and arm k, a state s_t^k in some compact state space S is realized, and arm k yields some bounded payoff $X^k(s_t^k)\in[-M,M]$, denoted X_t^k.

In each period t, the DM can choose exactly one out of the K arms, observe its state realization, and receive payoff from this arm. Let $a_t\in\{1,\ldots,K\}$ denote the arm chosen in period t, $s_t^{a_t}$ denote the state realization observed in period t, and $Z_t = X_t^{a_t}$ denote the payoff received in period t. Let $h_t=(a_1,s_1^{a_1},\ldots,a_t,s_t^{a_t})$ be the *partial history up to period* t, that is, the record of each arm chosen and state realization of the chosen arm in the first t periods. Thus h_t describes all information the DM has prior to choosing the period $t+1$ arm. $H_t=(\{1,\ldots,K\}\times S)^t$ is the set of all partial histories up to time t. In particular, $H_0=\emptyset$. A *strategy profile* is a random vector $\mathbf{a}=(a_1,\cdots,a_T)$, where $a_t: H_{t-1}\to\{1,\cdots,K\}$. So in the beginning of period $t+1$, the DM chooses which arm to pull next based on past history.

We incorporate ambiguity by allowing the DM to have multiple beliefs about the distribution governing the $\{s_t^k\}_{t=1}^T$ process. In particular, we adopt Marinacci [2002]'s model of multiple Bayesian priors. For fixed k, let Θ_k be a compact subset of $\Delta(S)$.[6] Conditional on θ_k, $\{s_t^k\}_{t=1}^T$ are i.i.d. with (unambiguous) likelihood distribution $l(\cdot|\theta_k)$. The DM's *a priori* information about the θ_k is modeled by a compact and convex set C_k of Bayesian priors on $\Delta(\Theta_k)$. When C_k is a singleton, there is no *a priori* ambiguity about arm k and $\{s_t^k\}_{t=1}^T$ reduces to a standard sampling process. We also assume that K arms are independent, so all prior information can be described by vector $C=(C_1,\ldots,C_K)$. We will refer to such a bandit problem as a (C,T)-bandit.

Finally we describe the *law of motion of beliefs*. At each partial history h_t, the DM is endowed with a vector of convex and compact posterior belief sets $C(\cdot|h_t)=(C_1(\cdot|h_t),\ldots,C_K(\cdot|h_t))$. Suppose in period $t+1$ the DM selects arm k and observes a state realization s_{t+1}^k. We assume her posterior belief sets conditional on history $h_{t+1}=(h_t,k,s_{t+1}^k)$ are updated in the following way:

- For arm k, belief set is updated at s_{t+1}^k prior-by-prior according to Bayes rule, so

$$C_k(\cdot|h_{t+1})=\{\mu_k(\cdot|s_{t+1}^k):\forall\mu_k\in C_k(\cdot|h_t)\}$$

- For any arm $j\neq k$, no updating occurs, so

$$C_j(\cdot|h_{t+1})=C_j(\cdot|h_t)$$

Note that the updated posterior belief sets $\{C_i(\cdot|h_{t+1})\}_{i=1}^K$ are also convex and compact. Since arms are *independent*, beliefs on arm $j\neq k$ are not updated at observation of s_{t+1}^k. Finally, we let $C(\cdot|h_{t+1})=(C_1(\cdot|h_{t+1}),\ldots,C_K(\cdot|h_{t+1}))$.

[5] See, for example, Berry and Fristedt [1985] for a textbook reference.

[6] For any separable metric space X, we used $\Delta(X)$ to denote the space of Borel probability measures on X.

3.2.2 Utility and Value Functions

Now we specify the utility. Suppose a DM follows a strategy $\mathbf{a}$ in a (C,T)-bandit and receives a random payoff stream $Z = (Z_1, \cdots, Z_T)$, where $Z_t = X_t^{a_t}$ for all t. In classic expected utility (EU) model (when $C = \{\mu\}$), his utility at each history h_t is simply the sum of discounted expected utility $U_{h_t}(\mu, T, \mathbf{a}) = E_{\mu(\cdot|h_t)}[\sum_{t'=t+1}^{T} \delta^{t'-t} Z_{t'}]$. In the multiple-prior case, we would like to consider a DM who adopts a maxmin expected utility (MEU, Gilboa and Schmeidler 1989, Epstein and Schneider 2003) criterion: he evaluates a random outcome by computing the EU with respect to every prior and considering the worst one. In a dynamic (C,T)-bandit problem, there are two obvious approaches to do this. One is to simply let $U'_{h_t}(C,T,\mathbf{a}) = \inf_{\mu \in C(\cdot|h_t)} E_\mu[\sum_{t'=t+1}^{T} \delta^{t'-t} Z_{t'}]$ for all h_t. Alternatively, $U_{h_t}(C,T,\mathbf{a})$ can be calculated backward recursively, applying maxmin EU criterion for the sum of instantaneous utility (Z_{t+1}) and discounted next-period continuation utility $(\delta U_{h_{t+1}}(C,T,\mathbf{a}))$ period-by-period for all h_t. These two approaches coincide in the special case of single-prior EU model, but differ in general whenever there is non-trivial ambiguity (C is not singleton). We will elaborate on this point in Section 3.5. For characterization of the optimal strategies in (C,T)-bandit problem, we take the second approach so that utilities are recursive and standard dynamic programming techniques apply.

Finite Horizon

For a fixed strategy $\mathbf{a}$, not all partial histories are consistent with $\mathbf{a}$. For example, if $\mathbf{a}(\emptyset) = 1$, then any partial history that starts with choosing arm 2 can never be reached under strategy $\mathbf{a}$. We say a partial history $h_t = (a'_0, s_1^{a'_0}, a'_1, \ldots, a'_t, s_t^{a'_t})$ is *consistent with strategy* $\mathbf{a}$ if $a_s(h_s) = a'_{s+1}$ for all $0 \le s < t$ and $h_s = (a'_0, s_1^{a'_0}, a'_1, \ldots, a'_s, s_s^{a'_s})$. For any strategy $\mathbf{a}$ and partial history h_t consistent with $\mathbf{a}$, we define ${}^{(h_t)}\mathbf{a}(\cdot) = \mathbf{a}(h_t, \cdot)$ to be $\mathbf{a}$'s continuation strategy on H_s/H_t, $s \ge t$.

For finite horizon problem $(T < \infty)$, we assume that the DM's *recursive multiple-priors utility* from following strategy $\mathbf{a}$ in a (C,T)-bandit can be computed backward recursively as follows: for all $\mathbf{a}$ and C,

1. the h_T-conditional utility $U_{h_T}(C,T,\mathbf{a}) = 0$ for any final history h_T consistent with strategy $\mathbf{a}$;

2. at any partial history h_t $(0 \le t < T)$ consistent with $\mathbf{a}$, her h_t-conditional utility from strategy $\mathbf{a}$ is

$$U_{h_t}(C,T,\mathbf{a}) = \inf_{\mu \in C(\cdot|h_t)} E_\mu[Z_{t+1} + \delta U_{h_{t+1}}(C,T,\mathbf{a})] \tag{3.1}$$

   where $h_{t+1} := (h_t, a(h_t), s_{t+1}^{a(h_t)})$ is the $t+1$-history follows from h_t, strategy $\mathbf{a}$, and random realization of state $s_{t+1}^{a(h_t)}$.

For simplicity, we use $U(C, T, \mathbf{a})$ to denote the ex-ante utility $U_{h_0}(C, T, \mathbf{a})$. We have two remarks on equation (3.1). First, we interpret payoff Z_t in terms of utils instead of consumptions or monetary rewards, a simplification that allows us to directly compute Z_{t+1} instead of $u(Z_{t+1})$ with some vNM risk index u. Second, conditional on h_t, $Z_{t+1} + \delta U_{h_{t+1}}(C, T, \mathbf{a})$ is a function of state $s_{t+1}^{a(h_t)}$ while μ is a product probability measure on $\prod_1^K \Delta(\Theta_k)$. When there is no confusion, for any $f : S \to \mathbb{R}$, we will simplify notation by using $E_{\mu_k}[f]$ for $\int_{\Theta_k} \int_S f(s^k) dl(s^k|\theta_k) d\mu_k(\theta_k)$.

The h_t-conditional value of a (C, T)-bandit problem is defined as the supremum of h_t-conditional utilities from following all strategies $\mathbf{a}$ that h_t is consistent with, that is,

$$V_{h_t}(C, T) = \sup_{\mathbf{a}} U_{h_t}(C, T, \mathbf{a})$$

In particular, the *value function* of a (C, T)-bandit problem is $V(C, T) = \sup_{\mathbf{a}} U(C, T, \mathbf{a})$. We say a strategy profile $\mathbf{a}^*$ is an *optimal strategy* if value $V(C, T)$ attains at $\mathbf{a}^*$, i.e., $V(C, T) = U(C, T, \mathbf{a}^*)$.

By the recursive construction of the utility function, applying standard dynamic programming techniques, we show that the value functions satisfy recursive relation (3.2) and an optimal strategy exists.

Proposition 3.1. *For any (C, T)-bandit, the conditional value functions satisfy the following recursive relation*

$$V_{h_t}(C, T) = \max_{a(h_t)=k} \inf_{\mu_k \in C_k(\cdot|h_t)} E_{\mu_k}[X_{t+1}^k + \delta \cdot V_{(h_t, k, s_{t+1}^k)}(C, T)] \tag{3.2}$$

Furthermore, there exists an optimal strategy $\mathbf{a}^*$ such that at all partial history h_t, the optimal choice is given by

$$a^*(h_t) \in \arg\max_k \inf_{\mu_k \in C_k(\cdot|h_t)} E_{\mu_k}[X_{t+1}^k + \delta \cdot V_{(h_t, k, s_{t+1}^k)}(C, T)]$$

This leads to the following corollary.

Corollary 3. *For any (C, T)-bandit, it is optimal to choose arm k at partial history h_t if and only if it is initially optimal to choose arm k in a $(C(\cdot|h_t), T - t)$-bandit.*

According to this corollary, to solve for full contingent optimal strategy profile for all bandits, it suffices to solve for what is initially optimal ($\mathbf{a}^*(\emptyset)$) for all bandits. This simplifies the characterization of optimal strategy.

Infinite Horizon

Next we extend the recursive construction of utilities to the infinite horizon case. In a (C, ∞)-bandit, a strategy profile $\mathbf{a}$ yields an infinite payoff stream $Z = (Z_1, \ldots, Z_t, \ldots)$. We construct the *infinite-horizon recursive multiple-priors utilities* $\{U_{h_t}(C, \mathbf{a})\}_{h_t}$ in the following way. For every finite T and partial history h_t ($t \leq T$), we define $U^T_{h_t}(C, \mathbf{a})$ to be the T-*truncated* version of $U_{h_t}(C, \mathbf{a})$, that is, the utility from following strategy $\mathbf{a}$ in a (C, T)-bandit, constructed backward recursively as in the finite horizon case. Then we set $U_{h_t}(C, \mathbf{a}) = \lim_{T\to\infty} U^T_{h_t}(C, \mathbf{a})$. The sequence $\{U^T_{h_t}(C, \mathbf{a})\}^\infty_{T=t+1}$ converges by Lemma C.1 in Appendix C.

Similarly we let the conditional value functions $V_{h_t}(C)$ be $\sup_{\mathbf{a}} U_{h_t}(C, \mathbf{a})$, and define the optimal strategy $\mathbf{a}^*$ as the strategy at which the value attains: $V(C) = U(C, \mathbf{a}^*)$. The next proposition shows that Proposition 3.1 for the finite-horizon case also holds for the infinite-horizon case.

Proposition 3.2. *For any (C, ∞)-bandit, the infinite-horizon recursive multiple-priors utilities $\{U_{h_t}(C, \mathbf{a})\}$ also satisfy recursive equation:*

$$U_{h_t}(C, \mathbf{a}) = \inf_{\mu\in C(\cdot|h_t)} E_\mu[Z_{t+1} + \delta U_{(h_t, a(h_t), s^{a(h_t)}_{t+1})}(C, \mathbf{a})]$$

The value functions $\{V_{h_t}(C)\}$ satisfy recursive equation:

$$V_{h_t}(C) = \max_{a(h_t)=k} \inf_{\mu_k\in C_k(\cdot|h_t)} E_{\mu_k}[X^k_{t+1} + \delta \cdot V_{(h_t,k,s^k_{t+1})}(C)]$$

There exists an optimal strategy $\mathbf{a}^*$ such that at all partial history h_t, the optimal choice is given by

$$a^*(h_t) \in \arg\max_k \inf_{\mu_k\in C_k(\cdot|h_t)} E_{\mu_k}[X^k_{t+1} + \delta \cdot V_{(h_t,k,s^k_{t+1})}(C)]$$

For all (C, ∞)-bandit, let $V^T_{h_t}$ be the value of its corresponding T-truncated finite horizon problem (C, T), that is, $V^T_{h_t}(C) := \sup_{\mathbf{a}} U^T_{h_t}(C, \mathbf{a})$. We show that the values of the infinite-horizon bandit problem can be approximated by the values of T-truncated finite horizon problems.

Proposition 3.3. *For all (C, ∞)-bandits and partial history h_t, $V_{h_t}(C) = \lim_{T\to\infty} V^T_{h_t}(C)$.*

3.2.3 One-armed Bandit

Next we look at the simple case when there are only two arms, and one of them is known and yields constant payoff. It captures the trade-off between experimenting with the unknown arm and exploiting the safe arm, and is traditionally called the one-armed bandit problem.[7]

Let us specify notations for a one-armed bandit problem (C, λ, T). Since only arm 1 is random, for every t the period state space is $S = S^1$ with σ-algebra $\mathcal{S}$, and the whole state space is $\Omega = S^T$ with σ-algebra $\Sigma = \sigma(\prod^T_1 \mathcal{S})$. Its natural filtration $\{\mathcal{F}_t\}^T_{t=1}$ is given by $\mathcal{F}_t := \sigma(s_1, \cdots, s_t)$ for all t. Any strategy profile $\mathbf{a} = (a_1, \cdots, a_T)$ is a predictable process,

[7]In some papers, this is also called two-armed bandit problem. To highlight the feature that there is only one arm with unknown distribution, we follow Berry and Fristedt [1985] and call it one-armed bandit problem.

as $a_t = \mathbf{a}(h_{t-1}) \in \mathcal{F}_{t-1}$. A random variable $N : \Omega \to \{0, 1\}$ is a *stopping time* if event $(N \leq t) \in \mathcal{F}_t$ for all t.

When C is a singleton set, say $\{\mu\}$, the problem reduces to a classic one-armed bandit problem. In this case, it is well known that in a (μ, λ, T)-bandit the optimal strategy is a stopping time, characterized by the *Gittins dynamic allocation index*[8]

$$\Lambda(\mu, T) := \max_{N \geq 1} \frac{E_\mu(\sum_{t=1}^N \delta^{t-1} X_t^1)}{E_\mu \sum_{t=1}^N \delta^{t-1}} \tag{3.3}$$

where N is a random stopping time.

The optimal strategy is a switching strategy characterized by the Gittins dynamic allocation index: start with arm 1 if and only if $\Lambda(\mu, T) \geq \lambda$, keep experimenting with arm 1 as long as $\Lambda(\mu(\cdot|s_1, \cdots, s_t), T - t) \geq \lambda$, switch to arm 2 the first time $\Lambda(\mu(\cdot|s_1, \cdots, s_t), T - t) < \lambda$, and stay with arm 2 until T.

For intuition about the optimal strategies in the general MP one-armed bandit problem (C, λ, T), we first look at a simple example with two periods and Bernoulli distributed arm one. In this case, the optimal strategy can be easily calculated by backward induction.

Example 3.1. Suppose $T = 2$, $K = 2$, and arm 2 yields constant payoff $\lambda = \frac{1}{2}$. Arm 1 is unknown. Let $S^1 = \{0, 1\}$, and s_t^1 has Bernoulli distribution with success rate $Pr(s_t^1 = 1|\theta) = \theta$ lying in $\Theta = [0, 1]$. Let arm 1's payoff be $X^1(s_t^1) = s_t^1$. Let $C_1 \subseteq \{\mu_\alpha = \alpha\delta_{\frac{3}{4}} + (1-\alpha)\delta_{\frac{1}{4}} | \alpha \in [0, 1]\}$ be the set of ex-ante beliefs about arm 1.[9] Discount rate is $\delta \in (0, 1)$.

Case 1: Suppose $C_1 = \{\mu_{\frac{1}{2}}\}$, then this reduces to a classic bandit problem. The optimal strategy is to choose arm 1 in period 1, continue to choose arm 1 in period 2 if $s_1^1 = 1$ (success), and switch to arm 2 in period 2 if $s_1^1 = 0$ (failure).

Case 2: Suppose $C_a = \{\mu_\alpha | \alpha \in [\frac{1}{2} - a, \frac{1}{2} + a]\}$, where $0 < a < \frac{1}{2}$ characterizes the degree of ambiguity. In this case, only two strategies are potentially optimal: (1) experiment with arm 1 in period 1, and continue to choose arm 1 in period 2 if and only if $s_1^1 = 1$ (success); (2) choose arm 2 in both periods. Since $C_a(\cdot|s_1^1 = 1) = [\frac{1.5-3a}{2-2a}, \frac{1.5+3a}{2+2a}]$ and $C_a(\cdot|s_1^1 = 0) = [\frac{0.5-a}{2+2a}, \frac{0.5+a}{2-2a}]$, the first strategy is optimal if and only if $V^{(1)} = \frac{1}{2} - a + \delta((\frac{1}{2} - a)\frac{1.5-3a}{2-2a} + (\frac{1}{2} + a) \cdot \frac{1}{2}) \geq V^{(2)} = \frac{1}{2}(1 + \delta)$. So the DM will start with experimenting with arm 1 if and only if $a \geq a^*(\delta)$, where $a^*(\delta) > 0$ is an increasing function of δ with value in $(0, 0.08)$.

This simple example generates three conjectures about optimal strategies in multiple-prior one-armed bandits.

1. The optimal strategy is still a switching strategy.
2. The more ambiguity averse a DM is (higher a), the less willing she is to experiment with the unknown arm.

[8] See, for example, Berry and Fristedt [1985] Chapter 5.

[9] δ_{θ_0} is the Dirac measure that assigns probability 1 to $(\theta = \theta_0)$.

3. For every multiple-prior set C and time horizon T, there exists an equivalent single-prior $\mu \in C$. That is, a multiple-prior DM with belief set C will behave just as an expected utility DM with the equivalent prior μ, for all levels of constant arm 2 payoff.

Conjecture 1 is formalized in Li [2012]. Below we state it without proof.

Theorem (Theorem 1, Li 2012). *For all (C, λ, T)-bandit, there exists a unique MP Gittins index, $\Lambda(C, T)$, such that*

- *the optimal strategy takes the form of a switching strategy: first experiment with arm 1 from period 1 to period N^*, then switch to arm 2 during period $N^* + 1$ to T.*
- *where N^* is a stopping time characterized by the first time the dynamic MP Gittins index falls below arm 2 payoff, i.e. $N^* := \inf\{t \geq 0 | \Lambda(C(\cdot|s_1^1, \cdots, s_t^1), T - t) < \lambda\}$.*

Moreover, the value of an (C, λ, T)-bandit is

$$V(C, \lambda, T) = \frac{1 - \delta^T}{1 - \delta} \max\{\Lambda(C, T), \lambda\} \tag{3.4}$$

Remark 4. Similar to the classic single-prior Gittins index, the multiple-prior Gittins index is characterized by the unique cutoff value of arm 2 payoff, at which the DM is indifferent between a strategy of always choosing arm 2 and a strategy of initially choosing arm 1 and continuing optimally. This is largely due to the backward recursive construction of the utilities, which ensures that introducing MP will not affect the tractable recursive structure of the classic bandit problem. However, this backward construction also comes at a cost. Unlike the classic Gittins index, which has a clean expression of $\Lambda(\mu, T) := \max_{N \geq 1} \frac{E_\mu(\sum_{t=1}^N \delta^{t-1} X_t^1)}{E_\mu \sum_{t=1}^N \delta^{t-1}}$, the MP Gittins index does not have a closed form solution. This is due to non-linearity of the maxmin EU operator. In Section 3.4, we will relate the MP Gittins index $\Lambda(C, T)$ to lower envelope of classic Gittins indices $\Lambda(\mu, T)$ for every prior μ in set C. It turns out that $\inf_{\mu \in C} \Lambda(\mu, T)$ is an upper bound for $\Lambda(C, T)$.

3.3 Comparative Statics

In this section, we examine the second conjecture from Example 3.1: increasing ambiguity (aversion) decreases the incentive to experiment with the unknown arm. This is true in our model.

We first need a general definition of "more ambiguity (averse) than". Following Ghirardato and Marinacci [2002], we say *i is more ambiguity (averse) than j* about the unknown arm if $C_j \subseteq C_i$. This can describe two sources of variations. First, i and j can be two different DMs, who face the same one-armed bandit problem and are given the same prior information, and DM i is more ambiguity averse than DM j. Second, it can be a single DM who faces two different one-armed bandit problems, where the unknown arm in problem i is provided with less precise prior information and thus more ambiguous than that in problem j. The comparative static statements in the next proposition hold for both interpretations.

Proposition 3.4 (Comparative Statics)**.** *For fixed λ and T, and one-armed bandits (C^i, λ, T) and (C^j, λ, T), if i is more ambiguity (averse) about arm 1 than j, i.e., $C_j \subseteq C_i$, then*

1. *i has a lower MP Gittins index than j, i.e., $\Lambda(C_i, T) \leq \Lambda(C_j, T)$.*
2. *i will experiment less than j, i.e., $N_i^* \leq N_j^*$.*

Proof. For part (1), by construction of utility

$$\begin{aligned} & C_j \subseteq C_i \\ & \Rightarrow U(C_i, \lambda, T, N) \leq U(C_j, \lambda, T, N), \forall N \\ & \Rightarrow \Delta(C_i, \lambda, T, N) \leq \Delta(C_i, \lambda, T, N), \forall N \\ & \Rightarrow \max_{N \geq 1} \Delta(C_i, \lambda, T, N) \leq \max_{N \geq 1} \Delta(C_i, \lambda, T, N) \end{aligned} \tag{3.5}$$

Since for all $k \in \{i, j\}$, $\Lambda(C_k, T)$ is the value of λ at which $\max_{N \geq 1} \Delta(C_k, \lambda, T, N) = 0$. By (3.5) and the fact that $\max_{N \geq 1} \Delta(C_k, \lambda, T, N)$ is strictly decreasing in λ, $\Lambda(C_i, T) \leq \Lambda(C_j, T)$ follows.

For part (2), $N_i^* \leq N_j^* \Leftrightarrow [\forall t, (N_i^* > t) \Rightarrow (N_j^* > t)] \Leftrightarrow [\Lambda(C_i(\cdot|h_t), T)) \geq \lambda \Rightarrow \Lambda(C_j(\cdot|h_t), T)) \geq \lambda]$. Since $C_j(\cdot|h_t) \subseteq C_i(\cdot|h_t)$ for all h_t, part (1) implies part (2). ☐

This generates interesting comparison between the effects of a change in ambiguity (aversion) and that of risk (aversion) on the optimal level of experimentation. Increasing ambiguity aversion and risk aversion both lead to lower level of experimentation. More surprisingly, for a given DM (so a fixed degree of ambiguity aversion/risk aversion), increasing the amount of risk and ambiguity in the problem have opposite effects on optimal experimentation. Under risk neutrality,[10] the value function is a convex function of arm 1 payoffs (X^1), so increasing risk (a mean-preserving spread of X^1) increases the value of experimentation. On the other hand, increasing ambiguity (an expansion of set C) decreases the minimum expected utility for any consumption process generated by experimentation, and leads to a lower incentive to experiment. This is illustrated by the following example.

Example 3.2 (Normal likelihood, normal priors, known variance)**.** Assume arm 1's payoff $X_t^1 = s_t^1$, and X_t^1 has normal likelihood $X_t^1|\theta \sim \mathcal{N}(\theta, \sigma^2)$ and normal conjugate priors $\theta|\tau, \sigma^2 \sim \mathcal{N}(\tau, \sigma^2)$. Let σ^2 be known and given. Denote a conjugate prior with mean τ and variance σ^2 by μ_{τ,σ^2}. We introduce ambiguity by allowing mean of the conjugate prior to take value in a closed real interval $[a, b]$, i.e., $C_{a,b,\sigma^2} = \{\mu_{\tau,\sigma^2} | \tau \in [a, b]\}$. Assume the safe arm 2 yields a constant payoff λ in every period.

First we state a fact about the value function of a unique prior $(\mu_{\tau,\sigma^2}, \lambda, T)$-bandit:[11]

[10] The claim here should not rely on the assumption of risk neutrality. A more careful definition of "increasing risk" should generalize the comparative static to allow for arbitrary risk attitude. This will be fixed later.

[11] Proof see Appendix C.

Fact 1. *For any given λ, T, $V(\mu_{\tau,\sigma^2}, \lambda, T)$ is weakly increasing in both τ and σ^2, and it is a convex function of τ.*

For the normal distribution case, we have[12]

$$V(C_{a,b,\sigma^2}, \lambda, T) = V(\mu_{a,\sigma^2}, \lambda, T) \tag{3.6}$$

$$\Lambda(C_{a,b,\sigma^2}, T) = \Lambda(\mu_{a,\sigma^2}, T) \tag{3.7}$$

Thus the comparative statics in ambiguity and risk are opposite:

1. Suppose there is an increase in ambiguity, say the prior set expands from $[a, b]$ to some bigger interval $[a', b']$, then by Fact 1, $V(\mu_{a',\sigma^2}, \lambda, T) \leq V(\mu_{a,\sigma^2}, \lambda, T))$, and applying (3.6) and (3.7)
$$\Lambda(C_{a,b,\sigma^2}, T) \geq \Lambda(C_{a',b',\sigma^2}, T)$$
Thus the DM is less likely to experiment with arm 1.

2. Suppose there is an increase in risk, say the variance expands from σ^2 to a higher $\tilde{\sigma}^2$, then by Fact 1 $V(\mu_{a,\sigma^2}, \lambda, T) \leq V(\mu_{a,\tilde{\sigma}^2}, \lambda, T)$, and applying (3.6) and (3.7)
$$\Lambda(C_{a,b,\sigma^2}, T) \leq \Lambda(C_{a,b,\tilde{\sigma}^2}, T)$$
Thus the DM is more likely to experiment with arm 1.

This suggests that ambiguity, instead of risk, may serve as an explanation to the widely observed underexperimentation of new products or technologies. In one-armed bandit problems, ambiguity generates testable behavioral implication that is *qualitatively* different from risk: within a given DM (and thus the degree of ambiguity aversion is fixed), increasing uncertainty (ambiguity) in the unknown arm decreases the incentive to experiment while increasing risk, by raising the option value of the unknown arm, increases the incentive to experiment.

3.4 Existence of Equivalent Prior

Here we want to examine the third conjecture: for every C and T, there exists an equivalent single-prior $\mu \in C$ such that a multiple-prior DM with belief set C will behave just as an expected-utility DM with the equivalent prior μ, in all (C, λ, T)-bandits. By Theorem 3.2.3, $\underline{\mu}$ is the *equivalent prior* for (C, λ, T)-bandit if $\underline{\mu} \in C$ and $\Lambda(C, T) = \Lambda(\underline{\mu}, T)$. So the question of whether an equivalence prior exists is equivalent to whether the multiple-prior Gittins index $\Lambda(C, T)$ is equal to the lower envelope of the single-prior Gittins indices, $\inf_{\mu \in C} \Lambda(\mu, T)$.

First we show that the multiple-prior Gittins index $\Lambda(C, T)$ is bounded above by the lower envelope of the classic single-prior Gittins indices.

[12] Proof see Appendix C.

Corollary 4. *For all one-armed bandits (C, λ, T),*

$$\Lambda(C, T) \leq \inf_{\mu \in C} \Lambda(\mu, T) \tag{3.8}$$

Proof. For any $\mu \in C$, let $C_j := \{\mu\}$, and $C_i := C$. By Proposition 3.4, $\Lambda(C, T) \leq \Lambda(\mu, T)$. Since this holds for every μ in C, $\Lambda(C, T) \leq \inf_{\mu \in C} \Lambda(\mu, T)$. □

Next we ask whether (3.8) holds with equality. If it is true, an equivalent prior exists for every one-armed bandit. Corollary 5 says that it also implies a minimax duality result: the utility of a multiple-prior DM who first chooses for each strategy the expected utility minimizing prior then the utility maximizing strategy, is equal to the utility from first choosing for each prior the expected utility maximizing strategy then the utility minimizing prior.[13] In many applications, the minimax equality serves as a trick to simplify characterization of optimal solutions. In our one-armed bandit problems, it can simplify the task of finding the multiple-prior Gittins index, an implicit solution, into computing the lower envelope of a set of single-prior Gittins index, an explicit expression given by (3.3).

Corollary 5. *The following statements are equivalent:*

1. *For all C and T, $\Lambda(C, T) = \inf_{\mu \in C} \Lambda(\mu, T)$.*
2. *For all C, T, and λ, $\sup_N U(C, \lambda, N, T) = \inf_{\mu \in C} \sup_N U(\mu, \lambda, N, T)$.*

Somewhat surprisingly, the equivalent prior may not always exist in one-armed bandit problems. Below is a counter example, when $\Lambda(C, T) < \inf_{\mu \in C} \Lambda(\mu, T)$. And as a result, for each prior μ in C and a μ-expected-utility DM, it is optimal to keep experimenting; while for a multiple-priors DM with prior set C, it is optimal to stop experimentation.

Example 3.3. Think of the case of Bernoulli likelihood. Now the DM considers three success rates possible: $\Theta_1 = \{0.1, 0.5, 0.9\}$. A prior will take the form of $\mu_a = a_1\delta_{0.1} + a_2\delta_{0.5} + a_3\delta_{0.9}$, for $a \in \mathbb{R}^3_+, a_1 + a_2 + a_3 = 1$. Consider prior set $C = \{\mu_a | (a_1 - \frac{1}{3})^2 + (a_2 - \frac{1}{3})^2 + (a_3 - \frac{1}{3})^2 \leq 0.16\}$, a 0.4-ball around the equiprobability prior in Δ^3. Let $T = 3$ and discount rate $\delta = 1$. Suppose $X^1_t = 100 \cdot 1_{\{s^1_t = 1\}}$, and arm 2 gives constant payoff $\lambda = 33$.

In (C, λ, T)-bandit, the DM's utility from choosing arm 1 initially and continuing optimally is strictly lower than the utility from switching to arm 2 at the very beginning, that is,

$$V^{(1)}(C, \lambda, T) = 98.67 < 99 = V^{(2)}(C, \lambda, T)$$

So it is optimal for the DM to switch to arm 2 at the very beginning in (C, λ, T)-bandit.

[13] This type of minimax duality theorems can be found in many previous applications of the multiple-prior utility. See, for example, Theorem 2 in Riedel [2009] for optimal stopping problem and Section 3.2 in Epstein and Wang [1994] for optimal portfolio choice.

Alternatively, for all $\mu \in C$, and the corresponding unique-prior bandit (μ, λ, T), the DM's utility from choosing arm 1 initially and continuing optimally is strictly higher than utility from switching to arm 2 at the very beginning, that is,

$$\inf_{\mu \in C} V^{(1)}(\mu, \lambda, T) = 100.36 > 99 = V^{(2)}(\mu, \lambda, T)$$

So for all $\mu \in C$, it is optimal for the DM to experiment with arm 1 for at least one period in (μ, λ, T)-bandit.

Thus in this example $\Lambda(C, T) < \inf_{\mu \in C} \Lambda(\mu, T)$.

The gap between $\Lambda(C, T)$ and $\inf_{\mu \in C} \Lambda(\mu, T)$ highlights an interesting dynamic aspect of how ambiguity (aversion) could affect optimal experimentation. The comparative statics in Proposition 3.4 describe the effect of ambiguity (aversion) on the incentive to experiment *at a give history node* h_t: having a set of conditional posterior $C(\cdot|h_t)$ instead of a unique posterior $\mu(\cdot|h_t)$ (in $C(\cdot|h_t)$) decreases the worst case evaluation of experimenting with arm 1, and thus lowers the incentive to experiment. In Example 3.3, when a multiple-prior DM with belief set C evaluates a strategy of experimenting with arm 1, she has the flexibility to use different minimizing priors for evaluation of arm 1 *at different history nodes*. Yet when an expected-utility DM with a single prior μ evaluates the same strategy, she is committed to using conditional posteriors from the same μ. This flexibility in choosing different minimizing prior μ at different history nodes, together with ambiguity aversion, makes experimentation even less attractive. So an ambiguity averse DM might experiment strictly less than any expected-utility DM with some prior μ within her prior set.

An interesting way to understand the difference is to view the MP bandit problem as a social experimentation under unanimity voting rule.[14] Consider the prior set C as a society of EU agents, each $\mu \in C$ representing an agent with belief μ about the payoff prospect from experimenting with some social reform (arm 1). If at each history node the society decides whether to experiment with arm 1 through unanimity voting rule, then the level of experimentation may be strictly less than, if the society votes through unanimity rule a dictator in period 0 and let this dictator decide whether to experiment at all subsequence history nodes. This is because in the former case, the pivotal voter can be different at different history nodes.

It is natural to ask whether we can characterize conditions under which $\Lambda(C, T) = \inf_{\mu \in C} \Lambda(\mu, T)$, and thus an equivalent prior exists. Here we look at the simple case when the prior set (C) and likelihood distributions (Θ) can be parameterized by one-dimensional real numbers, and generalize properties in Example 3.1 (two-point supported Bernoulli distribution) and Example 3.2 (normal-normal distribution). It turns out that the condition needed is quite strong.

[14] I thank David Ahn for suggesting this story. See Strulovici [2010] for an application of one-armed bandit problem on collective experimentation decisions under different voting rules.

Definition 3.1. For real intervals A, B, and function $f : A \times B \to \mathbb{R}_+$, let $f(\cdot|b) : A \to \mathbb{R}_+$ be a density function for every $b \in B$.[15] We say f has *monotone likelihood ratio property (MLRP)* in a if for all $a_1 < a_2$ and $b_1 < b_2$,

$$\frac{f(a_1|b_2)}{f(a_1|b_1)} \leq \frac{f(a_2|b_2)}{f(a_2|b_1)}.$$

In particular, MLRP implies that for $b_1 < b_2$, $f(\cdot|b_2)$ first order stochastically dominates (FOSD) $f(\cdot|b_1)$. MLRP, instead of FOSD or other stochastic orders, serves our purpose because it implies FOSD and can be preserved after Bayesian updating.[16]

Condition 1. *Suppose both Θ and C can be parameterized by one dimensional real intervals and density functions exist. Let $\Theta = [\underline{\theta}, \bar{\theta}]$ and $C = \{\mu_a \in \Delta(\Theta) : a \in [\underline{a}, \bar{a}]\}$. Furthermore, suppose $\{l(s|\theta)\}_{\theta\in\Theta}$ has MLRP in s, and $\{\mu_a \in \Delta(\Theta) : a \in [\underline{a}, \bar{a}]\}$ has MLRP in θ.*

Proposition 3.5 (Existence of Equivalent Prior)**.** *For one-armed bandit (C, λ, T), if C and Θ satisfy Condition 1, then $\Lambda(C, T) = \Lambda(\mu_{\underline{a}}, T)$. In this case, $\mu_{\underline{a}}$ is the equivalent prior for (C, T).*

Proof. See Appendix. □

To sum up, in this section we highlight a discrepancy between the multiple-prior Gittins index and the lower envelope of classic single-prior Gittins indices for every prior in the belief set, and characterize a restrictive set of conditions under which this discrepancy disappears.

3.5 Discussion

In Section 3.2, we construct the multiple-prior utility from following a strategy backward recursively, applying maxmin EU criterion period by period. Alternatively, we could have applied maxmin EU criterion to the full discounted consumption stream and defined $U'_{h_t}(C, T, \mathbf{a}) = \inf_{\mu\in C(\cdot|h_t)} E_\mu[\sum_{t'=t+1}^{T} \delta^{t'-t} Z_{t'}]$ for all h_t. The next proposition says that these two specifications coincide on all consumption streams if and only if the set of predictive distribution of C, $\mathcal{P}_C = \{P_\mu = \int_\Theta l(\cdot|\theta)d\mu : \mu \in C\}$, is singleton. That is, when ambiguity disappears.

Proposition 3.6. *For any bounded adapted payoff process $(Z_1, \cdots, Z_T)$,*

$$U(C)(Z_1, \cdots, Z_T) \leq \inf_{\mu\in C} U(\{\mu\})(Z_1, \cdots, Z_T) \tag{3.9}$$

And the inequality holds strictly for some $(Z_1, \cdots, Z_T)$ if and only if $\mathcal{P}_C = \{\int_\Theta l(\cdot|\theta)d\mu : \mu \in C\}$ is non-singleton.

[15] $\int_A f(a|b)da = 1$ for all $b \in B$.

[16] This is proved in proof of Proposition 3.5. Other commonly seemed properties, like FOSD or single crossing properties, cannot be preserved after Bayesian updating and thus are insufficient for our purpose. See Klemens [2007] for a discussion for this.

Proof. See Appendix C. □

This says, in a repeated sampling setting, exchangeable probabilities (right hand side of (3.9)), recursive utility (left hand side of (3.9)), and non-trivial ambiguity cannot be satisfied at the same time. This modeling trade-off has been shown in generality in Epstein and Seo [2011].[17] Here we give a direct proof for the case of multiple-prior utility.[18] In the multiple-prior one-armed bandit problems, as studied in Section 3.3, this discrepancy generates the gap between the multiple-prior Gittins index and the lower envelop of all single-prior Gittins indices, and behaviorally, that a recursive multiple-prior-utility DM might have strictly less incentive to experiment than any expected-utility DM with belief lying in her prior set.

[17]It is well known that imposing dynamic consistency results in restrictions on the ambiguity representation. See, for example, Epstein and Schneider [2003], Maccheroni et al. [2006b], Siniscalchi [2011].

[18]Slightly different from our set-up, Epstein and Seo [2011] does not consider intermediate consumptions.

Bibliography

B. Acciaio, H. Follmer, and I. Penner. Risk assessment for uncertain cash flows: Model ambiguity, discounting ambiguity, and the role of bubbles. *Forthcoming in Finance and Stochastics*, 2011.

C. M. Anderson. Ambiguity aversion in multi-armed bandit problems. *Theory and Decision*, 72:15–33, 2012.

D. K. Backus, B. R. Routledge, and S. E. Zin. *Exotic Preferences for Macroeconomists*, volume 19 of *NBER Macroeconomics Annual 2004*, chapter 5, pages 319 – 414. MIT Press, 2005.

C. Bellemare, M. Krause, S. Kroger, and C. Zhang. Myopic loss aversion: Information feedback vs. investment flexibility. *Economics Letters*, 87:319–324, 2005.

D. Bergemann and J. Valimaki. *The New Palgrave Dictionary of Economics*, chapter Bandit Problems. 2008.

D. A. Berry and B. Fristedt. *Bandit Problems: Sequential Allocation of Experiments*. Chapman and Hall, 1st edition, 1985.

G. Bryan. Ambiguity and insurance. *Working paper*, 2010.

H. Cao, T. Wang, and H. Zhang. Model uncertainty, limited market participation, and asset prices. *Review of Financial Studies*, 18(4):1219–1251, 2005.

D. Dillenberger. Preferences for one-shot resolution of uncertainty and allais-type behavior. *Econometrica*, 78:1973–2004, 2010.

J. Eichberger, S. Grant, and D. Kelsey. Updating choquet beliefs. *Journal of Mathematical Economics*, 43:888–899, 2007.

L. Epstein and M. Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113: 1–31, 2003.

L. Epstein and T. Wang. Intertemporal asset pricing under knightian uncertainty. *Econometrica*, pages 283–322, 1994.

L. G. Epstein. A definition of uncertainty aversion. *Review of Economic Studies*, 66:579–608, 1999.

L. G. Epstein and M. LeBreton. Dynamically consistent beliefs must be bayesian. *Journal of Economic Theory*, 61:1–22, 1993.

L. G. Epstein and M. Schneider. Ambiguity and asset markets. *Annual Review of Financial Economics*, 2:315–346, 2010.

L. G. Epstein and K. Seo. Symmetry or dynamic consistency? *The B.E. Journal of Theoretical Economics*, 11, 2011.

L. G. Epstein and J. Zhang. Subjective probabilities on subjectively unambiguous events. *Econometrica*, 69:265–306, 2001.

L. G. Epstein and S. E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 1989.

L. G. Epstein and S. E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *The Journal of Political Economy*, 99(2):263–286, 1991.

P. Ghirardato and M. Marinacci. Ambiguity made precise: A comparative foundation. *Journal of Economic Theory*, 102:251–289, 2002.

I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

J. Gittins, K. Glazebrook, and R. Weber. *Multi-armed Bandit Allocation Indices*. John Wiley & Sons, Ltd, 2nd edition, 2011.

J. C. Gittins and D. M. Jones. A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, 66:561–565, 1979.

U. Gneezy, A. Kapteyn, and J. Potters. Evaluation periods and asset prices in a market experiment. *The Journal of Finance*, 52(2):821–837, 2003.

S. Grant, A. Kajii, and B. Polak. Intrinsic preference for informations. *Journal of Economic Theory*, 83:233–259, 1998.

S. Grant, A. Kajii, and B. Polak. Temporal resolution of uncertainty and recursive non-expected utility models. *Econometrica*, 68(2):425–434, 2000.

M. Haigh and J. List. Do professional traders exhibit myopic loss aversion? An experimental analysis. *The Journal of Finance*, 60:523–534, 2005.

Y. Halevy. Ellsberg revisited: An experimental study. *Econometrica*, 75(2):503–536, 2007.

E. Hanany and P. Klibanoff. Updating preferences with multiple priors. *Theoretical Economics*, 2:261–298, 2007.

E. Hanany and P. Klibanoff. Updating ambiguity averse preferences. *The B.E. Journal of Theoretical Economics*, 9, 2009.

L. P. Hansen and T. J. Sargent. Robust control and model uncertainty. *The American Economic Review , Papers and Proceedings*, 91:60–66, 2001.

L. P. Hansen and T. J. Sargent. *Robustness.* Princeton University Press, 2007.

I. Herstein and J. Milnor. An axiomatic approach to measurable utility. *Econometrica*, 21: 291–297, 1953.

N. Ju and J. Miao. Ambiguity, learning, and asset returns. *Working paper*, 2010.

B. E. Kahn and R. K. Sarin. Modeling ambiguity in decisions under uncertainty. *Journal of Consumer Research*, 15:265–272, 1988.

B. Klemens. When do ordered prior distributions induce ordered posterior distributions? *mimeo*, 2007.

P. Klibanoff, M. Marinacci, and S. Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73:1849–1892, 2005.

P. Klibanoff, M. Marinacci, and S. Mukerji. Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144:930–976, 2008.

B. Koszegi and M. Rabin. Reference-dependent consumption plans. *American Economic Review*, 99:909–936, 2009.

D. Kreps. *Notes on the Theory of Choice.* Westview Press, 1988.

D. M. Kreps and E. L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46:185–200, 1978.

J. Li. K-armed bandit with multiple-prios. *Working paper*, 2012.

F. Maccheroni, M. Marinacci, and A. Rustichini. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74:1447–1498, 2006a.

F. Maccheroni, M. Marinacci, and A. Rustichini. Dynamic variational preferences. *Journal of Economic Theory*, 128:4–44, 2006b.

M. Marinacci. Learning from ambiguous urns. *Statistical Papers*, 43:143–151, 2002.

J. Miao and N. Wang. Risk, uncertainty, and option exercise. *Journal of Economic Dynamics and Control*, 35:442–461, 2011.

R. F. Nau. Uncertainty aversion with second-order utilities and probabilities. *Management Science*, 52(1):136–145, 2006.

K. G. Nishimura and H. Ozaki. Search and knightian uncertainty. *Journal of Economic Theory*, 119:299–333, 2004.

C. P. Pires. A rule for updating ambiguous beliefs. *Theory and Decision*, 53:137–152, 2002.

F. Riedel. Dynamic coherent risk measures. *Stochastic Processes and their Applications*, 112:185 200, 2004.

F. Riedel. Optimal stopping with multiple priors. *Econometrica*, 77:857–908, 2009.

D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–87, 1989.

U. Segal. Two-stage lotteries without the reduction axiom. *Econometrica*, 58(2):349–377, 1990. ISSN 00129682.

K. Seo. Ambiguity and second-order belief. *Econometrica*, 77:1575–1605, 2009.

M. Siniscalchi. Dynamic choice under ambiguity. *Theoretical Economics*, 6(3):379–421, 2011.

C. Skiadas. Recursive utility and preferences for information. *Economic Theory*, 12:293–312, 1998.

B. Strulovici. Learning while voting: Determinants of collective experimentation. *Econometrica*, 78(3):933–971, 2010.

T. Strzalecki. Temporal resolution of uncertainty and recursive models of ambiguity aversion. *Working paper*, 2010.

T. Strzalecki. Axiomatic foundations of multiplier preferences. *Econometrica*, 79:47–73, 2011.

Appendix A

Proofs for Chapter 1

Lemma 1.2

Proof of Lemma 1.1. Fix (π, f). We want to show that the sets $U = \{(\pi', g) : (\pi', g) \succcurlyeq (\pi, f)\}$ and $L = \{(\pi', g) : (\pi, f) \succcurlyeq (\pi', g)\}$ are closed.

Let $\{(\pi'_n, g_n)\}$ be a convergent sequence in the set U, with limit (π', g). We want to show (π', g) is also in U. Since $\pi'_n \to \pi'$ in the discrete topology on Π, there exists some N such that for all $n > N$, $\pi'_n = \pi'$. Continuity of $\succcurlyeq_\pi$ ensures there exists a constant act x_f, with $(\pi, f) \sim (\pi, x_f) \sim (\pi', x_f)$, where the last statement follows from Stable Risk Preferences. If $(\pi', x_f) \sim (\pi, f) \succ (\pi', g)$, then by continuity of $\succcurlyeq_{\pi'}$, there exists $M (> N)$ such that for all $n > M$, $(\pi', x_f) \succ (\pi', g_n)$. So $(\pi, f) \succ (\pi'_n, g_n)$ for sufficiently large n, a contradiction to the assumption $\{(\pi'_n, g_n)\} \subseteq U$. □

The next lemma verifies the existence of certainty equivalents as result of Continuity and Monotonicity.

Lemma A.1. *For any nonempty $E \in \Sigma$, if $\succcurlyeq_E$ satisfies Continuity and Monotonicity, then for every act f we can find an E-conditional certainty equivalent $c(f|E) \in X$ such that $c(f|E) \sim_E f$.*

Proof. Let $f \in \mathcal{F}$. Since f is finitely ranged, by Monotonicity there exists $x^*, x_* \in X$ such that $x^* \succcurlyeq_E f \succcurlyeq_E x_*$. By continuity, $U = \{\alpha \in [0,1] : \alpha x^* + (1-\alpha)x_* \succcurlyeq_E f\}$ and $L = \{\alpha \in [0,1] : f \succcurlyeq_E \alpha x^* + (1-\alpha)x_*\}$ are closed subsets of $[0,1]$. Since $U \cup L = [0,1]$, by connectedness of $[0,1]$, $U \cap L \neq \emptyset$. Thus there exists $c(f|E) \in U \cap L$, and by definition $c(f|E) \sim_E f$. □

Proof of Lemma 1.2. First we show equivalence of the two statements. (2) $\Rightarrow$ (1) is a straightforward verification. We show (1) $\Rightarrow$ (2).

Since $\succcurlyeq$ is a continuous and monotone preference relation, there exists a continuous function $V : \Pi \times \mathcal{F} \to \mathbb{R}$ that represents $\succcurlyeq$. By Time Neutrality, $(\pi_0, f) \sim (\pi^*, f), \forall f$, so $V(\pi_0, \cdot) = V(\pi^*, \cdot) : \mathcal{F} \to \mathbb{R}$ represents the restricted preference relations $\succcurlyeq_{\pi_0}$ and $\succcurlyeq_{\pi^*}$. Let $V_0(\cdot) := V(\pi_0, \cdot)$, and let $u : X \to \mathbb{R}$ be the restriction of V_0 to constant acts X, where

$u(x) = V_0(x)$. Since V_0 is continuous, u is continuous. $u(X)$ is connected and thus an interval in $\mathbb{R}$, since $X = \Delta(Z)$ is connected. We define the functional $I_0 : B(\Sigma, u(X)) \to \mathbb{R}$ by $I_0(\xi) = V_0(f)$, where $\xi \in B(\Sigma, u(X))$, $f \in \mathcal{F}$ satisfy $u \circ f = \xi$. Then I_0 is well-defined and monotone by monotonicity of $\succsim_0$. For any $k \in u(X)$, choose the constant act x such that $u(x) = k$. Then by definition, $I_0(\bar{k}) = V_0(x) = u(x) = k$. So I_0 is normalized.

Similarly, for every π, the continuous function $V_\pi = V(\pi, \cdot) : \mathcal{F} \to \mathbb{R}$ represents $\succsim_\pi$. We show the connection between V_π and (u, I_0). Fix $\pi, f, E_i \in \pi$. By continuity and monotonicity of $\succsim_{E_i}$, we can find conditional certainty equivalent $c(f|E_i) \sim_{E_i} f$. By π-recursivity, $(\pi, f) \sim (\pi, c(f|\pi))$. Then $(\pi, c(f|\pi)) \sim (\pi^*, c(f|\pi)) \sim (\pi_0, c(f|\pi))$, where the first indifference is by Indifference to Redundant Information, and the second by Time Neutrality. By transitivity of $\succsim$, $(\pi, f) \sim (\pi_0, c(f|\pi))$, so $V_\pi(f) = V_0(c(f|\pi)) = I_0(u \circ c(f|\pi))$. □

$\succsim_0$-non-null Events

This subsection clarifies the concept of a $\succsim_0$-non-null event for defining conditional preferences.

The literature normally adopts the condition of a non-null event from Savage. An event E is *Savage $\succsim_0$-non-null* if there exists f, g, h, such that $fEh \succ_0 gEh$.

We consider a stronger condition: an event E is *$\succsim_0$-non-null* if there exist constant acts x^*, x_* such that $x^* \succ_0 x_*$ and $x^*Ex_* \succ_0 x_*$. An event E is Savage $\succsim_0$-non-null if it is $\succsim_0$-non-null, but not vice versa. The next lemma compares how these two definitions differ in the variational preference family.

Lemma A.2. *Suppose $\succsim_0$ has a variational representation (u, c). An event E is $\succsim_0$-non-null if and only if $p(E) > 0$ for all $p \in c^{-1}(0)$. An event E is Savage $\succsim_0$-non-null if and only if there exists some act f and some $p \in \arg\min_{p' \in \Delta(S)} \int u(f)dp' + c(p')$ such that $p(E) > 0$.*

Proof. For the first claim, we prove E is $\succsim_0$-non-null iff $\exists p \in c^{-1}(0)$ such that $p(E) = 0$. Choose constant acts x^*, x_* such that $x^* \succ_0 x_*$. First, suppose $\exists p \in c^{-1}(0)$ such that $p(E) = 0$. Then

$$V_0(x^*Ex_*) = u(x^*)p(E) + u(x_*)p(E^c) + c(p) = u(x_*) = V_0(x_*)$$

The first equality holds because $c(p) = 0$ and $p(E) = 0$. Next, suppose instead $p(E) > 0$ for all $p \in c^{-1}(0)$. Then let $p^* \in \arg\min_{p'} u(x^*)p'(E) + u(x_*)p'(E^c) + c(p')$. Either $p^* \in c^{-1}(0)$ and $p^*(E) > 0$, or $c(p^*) > 0$. In either case,

$$V_0(x^*Ex_*)u(x^*)p^*(E) + u(x_*)p^*(E^c) + c(p^*) > u(x_*)$$

so $x^*Ex_* \succ_0 x^*$.

For the second claim, suppose there exists some act f and some $p \in argmin_{p' \in \Delta(S)} \int u(f)dp' + c(p')$ such that $p(E) > 0$. Then we can construct an act f' such that $f'(s) = f(s)$ for all

$s \in E^c$, and $u(f'_s) = u(f_s) - \epsilon$ for all $s \in E$, and some $\epsilon > 0$. Since $p(E) > 0$,

$$\begin{aligned} V_0(f) &= \int_E u(f)dp + \int_{E^c} u(f)dp + c(p) \\ &> \int_E u(f)dp - \epsilon p(E) + \int_{E^c} u(f)dp + c(p) \\ &= \int_S u(f')dp + c(p) \geq V_0(f') \end{aligned}$$

So $f \succ_0 f'$. For the converse, suppose there exists f, g, h such that $fEh \succ_0 gEh$. Let $p \in \arg\min_{p' \in \Delta(S)} \int u(gEh)dp' + c(p')$. We argue that $p(E) > 0$. If instead $p(E) = 0$, then

$$V_0(gEh) = \int_E u(g)dp + \int_{E^c} u(h)dp + c(p) = \int_E u(f)dp + \int_{E^c} u(h)dp + c(p) \geq V_0(fEh)$$

This contradicts $fEh \succ_0 gEh$. □

Suppose $\succcurlyeq_0$ has an MEU representation $(u, \mathcal{P})$. As a corollary, E is $\succcurlyeq_0$-non-null if and only if $p(E) > 0$ for all $p \in \mathcal{P}$. In contrast, E is Savage $\succcurlyeq_0$-non-null if and only if there exists f and $p \in \arg\min_{p \in \mathcal{P}} \int u(f)dp$ such that $p(E) > 0$.

For the results about updating, the stronger $\succcurlyeq_0$-non-null condition is needed. Pires [2002] shows that if the unconditional preferences $\succcurlyeq_0$ have an MEU representation $(u, \mathcal{P})$ and all priors give positive probability to event E, then Conditional Certainty Equivalence Consistency is satisfied if and only if $\succcurlyeq_E$ has an MEU representation $(u, \mathcal{P}_E)$, where $\mathcal{P}_E$ is the prior-by-prior updated posteriors from $\mathcal{P}$. In Section 1.4.3, we show that if the unconditional preferences $\succcurlyeq_0$ have a variational representation (u, c) and $p(E) > 0$ for all $p \in c^{-1}(0)$, then Conditional Certainty Equivalence Consistency is satisfied if and only if $\succcurlyeq_E$ has a variational representation (u, c_E), where c_E is obtained from c using update rule (A.1). In both cases, E has to be $\succcurlyeq_0$-non-null instead of Savage $\succcurlyeq_0$-non-null.

In the text, we impose Strong Monotonicity on $\succcurlyeq_0$ to ensure that updating is always well-defined. The following lemma follows directly by definition.

Lemma A.3. *If $\succcurlyeq_0$ satisfies Strong Monotonicity, then every event E in Σ is $\succcurlyeq_0$-non-null.*

Theorem 1.1

We first recall a result from Maccheroni et al. [2006a].

Lemma 28, Maccheroni et al. [2006a] *A binary relation $\succcurlyeq_0$ on $\mathcal{F}$ satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, and Non-degeneracy if and only if there exists a nonconstant affine function $u : X \to \mathbb{R}$ and a normalized, monotone, and translation invariant $I_0 : B(S, u(X)) \to \mathbb{R}$ such that*

$$f \succcurlyeq_0 g \Leftrightarrow I_0(u(f)) \geq I_0(u(g))$$

Below we will apply this result to prove our representation Theorem 1.1.

Proof of Theorem 1.1. We verify only the direction (1) $\Rightarrow$ (2). The other direction is straightforward.

By Lemma 1.2, (i) implies there exists a continuous function $V_0 : \mathcal{F} \to \mathbb{R}$ such that that for each π, $\succcurlyeq_\pi$ can be represented by

$$V(\pi, f) = V_0(c(f|\pi))$$

where $c(\cdot|\pi) : \mathcal{F} \to \mathcal{F}_\pi$ is the conditional certainty equivalent mapping.

Define $u : X \to \mathbb{R}$ by $u(x) = V_0(x)$. Define $I_0 : B(\Sigma, u(X)) \to \mathbb{R}$ by $I_0(\xi) = V_0(f)$, for $\xi \in B(\Sigma, u(X))$, $f \in \mathcal{F}$ such that $u \circ f = \xi$. By Lemma 28 in Maccheroni et al. [2006a], Weak Certainty Independence, Continuity, Monotonicity, and Non-degeneracy of $\succcurlyeq_0$ implies that u is continuous, nonconstant and affine, and I_0 is well-defined, continuous, normalized, and translation invariant. Moreover, for any $\xi, \xi' \in B(S, u(X))$ such that $\xi > \xi'$, there exists $f, g \in \mathcal{F}$ such that $u \circ f = \xi$ and $u \circ g = \xi'$, $f(s) \succcurlyeq_0 g(s)$ for all s, and $f(s) \succ_0 g(s)$ for some s. By Strong Monotonicity of $\succcurlyeq_0$, $f \succ_0 g$ and thus $I_0(\xi) > I_0(\xi')$. So I_0 is strongly monotone.

Next we show that for all f and nonempty $E \in \Sigma$, $k = I_0[(u \circ f)E\bar{k}]$ has a unique solution in $u(X)$.

Existence. Fix f and nonempty E. Define $G(k) = I_0[(u \circ f)E\bar{k}] - k = I_0[(u \circ f - \bar{k})E0]$, for all $k \in u(X)$. Since f is finite-ranged, we can find x^*, x_* such that $x^* \succ_0 f(s) \succ_0 x_*$ for all s. Let $k^* = u(x^*)$, and $k_* = u(x_*)$. Then $G(k^*) \geq 0$, and $G(k_*) \leq 0$ by monotonicity of I_0. Since I_0 is continuous, G is a continuous function of k on $u(X)$. By the intermediate value theorem, there exists $k_0 \in [k_*, k^*]$ such that $G(k_0) = 0$.

Uniqueness. Suppose k_1 and k_2 both solve $k = I_0[(u \circ f)E\bar{k}]$, and $k_1 \neq k_2$. Without loss of generality, let $k_1 > k_2$. By translation invariance of I_0,

$$I_0[(u \circ f - \bar{k}_1)E0] = I_0[u(fE\bar{k}_1)] - k_1 = 0 = I_0[u(fE\bar{k}_2)] - k_2 = I_0[(u \circ f - \bar{k}_2)E0]$$

Then $(u \circ f - \bar{k}_1)E0 < (u \circ f - \bar{k}_2)E0$, since E is non-empty. Since I_0 is strictly monotone, $I_0[(u \circ f - \bar{k}_1)E0] < I_0[(u \circ f - \bar{k}_2)E0]$. A contradiction.

For any $\pi = \{E_1, E_2, \cdots, E_n\}$, by Conditional Certainty Equivalent Consistency, x_i is the E_i-conditional Certainty Equivalent of f if and only if $x_i \sim_0 fE_ix_i$. This implies that $u(x_i)$ solves $k = I_0[(u \circ f)E_i\bar{k}]$ and $x_i \sim_0 c(f|E_i)$. So $u(x_i) = u(c(f|E_i))$, which implies $V_0(f|\pi) = u \circ c(f|\pi)$. As a result, $V(\pi, f) = V_0(c(f|\pi)) = I_0(V_0(f|\pi))$ by definition of I_0.

Finally we prove uniqueness. If both u and u' are affine representations of $\succcurlyeq_0$ on X, by the Mixture Space Theorem [Herstein and Milnor, 1953], $u' = au + b$ for some $a, b \in \mathbb{R}$, and $a > 0$. If both (u, I_0) and (u, I_0') represent $\succcurlyeq_0$, then there exists a strictly increasing $\phi : u(X) \to u(X)$ such that $I_0'(\xi) = \phi(I_0(\xi))$ for all $\xi \in B(S, u(X))$. Since I_0 and I_0' are normalized, for any $k \in u(X)$, $I_0'(\bar{k}) = k = \phi(I_0(\bar{k})) = \phi(k)$. This implies that ϕ is the identity mapping, so $I_0' = I_0$. $\square$

Theorem 1.2

Proof. By Conditional Certainty Equivalent Consistency, $fEx \sim_0 x \Leftrightarrow x \sim_E f$, for all f, x, E. So it suffices to show that $\succcurlyeq$ exhibits partial information aversion if and only if $x \sim_E f \Rightarrow f \succcurlyeq_0 xEf$, for all f, x, E.

Suppose $\succcurlyeq_0$ satisfies Event Complementarity. Fix a finite partition $\pi = \{E_1, \cdots, E_n\}$, and an act f. For each $i = 1, \cdots, n$, let $x_i \in X$ be the E_i-conditional certainty equivalent of f, i.e., $x_i \sim_{E_i} f$. Let $f_0 := f$, $f_1 = x_1 E_1 f_0$, $f_2 = x_2 E_2 f_1$, $\cdots$, $f_n = x_n E_n f_{n-1} = (x_1 E_1 x_2 E_2 \cdots x_{n-1} E_{n-1} x_n)$. Note that f_n is π-measurable. Also $x_i \sim_{E_i} f_{i-1}, \forall i = 1, \cdots, n$, thus $(\pi_0, f_{i-1}) \succcurlyeq (\pi_0, f_i)$ by Event Complementarity, and $(\pi, f_0) \sim (\pi, f_1) \sim \cdots \sim (\pi, f_n)$ by π-Recursivity. Putting these results together yields:

$$\begin{aligned}
(\pi, f) &\sim (\pi, f_n) \sim (\pi^*, f_n) \\
&\sim (\pi_0, f_n) \quad \text{(by Time Neutrality)} \\
&\preccurlyeq (\pi_0, f_{n-1}) \cdots \preccurlyeq (\pi_0, f)
\end{aligned}$$

Since this is true for an arbitrary act f and partition π, $\succcurlyeq$ exhibits aversion to partial information.

We prove the converse by contradiction. Suppose not, so $\succcurlyeq$ exhibits aversion to partial information but there exists some π, $E \in \pi$, f, and x such that $f \sim_E x$, but $(\pi_0, xEf) \succ (\pi_0, f)$. Let $n_1, \cdots, n_m$ be labels for states in E^c, i.e., $E^c = \{s_{n_1}, \cdots, s_{n_m}\}$. Then consider the finer partition $\pi' = \{E, \{s_{n_1}\}, \cdots, \{s_{n_m}\}\}$. Thus xEf is π'-measurable, and by Axioms 4 and 5, $(\pi', xEf) \sim (\pi^*, xEf) \sim (\pi_0, xEf)$. By π'-Recursivity, $(\pi', f) \sim (\pi', xEf)$. By transitivity, $(\pi', f) \sim (\pi_0, xEf) \succ (\pi_0, f)$. This violates partial information aversion, a contradiction.

□

Proposition 1.2

Proof. For part (1), by Theorem 1.2, it suffices to show that $\succcurlyeq_0$ satisfies Event Complementarity. Since $\succcurlyeq_0$ belongs to the MEU class, by Lemma 3.3 in Gilboa and Schmeidler [1989], I_0 is superadditive. Event Complementarity follows from that.

For part (2), if $\succcurlyeq_0$ has an MEU representation $(u, \mathcal{P})$ and $\succcurlyeq$ is recursively generated by $\succcurlyeq_0$, then $\succcurlyeq$ can be represented by

$$\begin{aligned}
V(\pi, f) &= \min_{p \in \mathcal{P}} \sum_{i=1}^{n} [\min_{p^i \in \mathcal{P}} \int u(f) dp^i(\cdot|E_i)] p(E_i) \\
&= \min_{p \in \mathcal{P}} \min_{p^i \in \mathcal{P}} \sum_{i=1}^{n} [\int u(f) dp^i(\cdot|E_i)] p(E_i) \\
&= \min_{p' \in rect_\pi(\mathcal{P})} \int u(f) dp'
\end{aligned}$$

Suppose $\mathcal{P}$ is not π-rectangular, so there exists $q \in rect_\pi(\mathcal{P})\backslash\mathcal{P}$. Since $\mathcal{P}$ is convex and compact, by the strict separating hyperplane theorem, there exists a nonzero, bounded and measurable map $\xi \in B(\Sigma, \mathbb{R})$ such that

$$\int \xi dq < \int \xi dp, \forall p \in \mathcal{P}$$

Without loss of generality, let $0 \in int(u(X))$. There exists $f \in \mathcal{F}$ such that $u(f) = \alpha\xi$, for some $\alpha > 0$. Thus without loss of generality we can replace ξ by $u(f)$ in above inequality. By compactness of $\mathcal{P}$, $\min_{p\in\mathcal{P}} \int u(f)dp$ attains at some $p^* \in \mathcal{P}$, so using above

$$V(\pi, f) = \min_{q' \in rect_\pi(\mathcal{P})} \int u(f)dq' \leq \int u(f)dq < \int u(f)dp^* = V(\pi_0, f)$$

Thus $\succcurlyeq$ is strictly averse to partition π at f.

For the converse, suppose $\mathcal{P}$ is π-rectangular, so $\mathcal{P} = rect_\pi(\mathcal{P})$. Then $V(\pi, f) = V(\pi_0, f), \forall f$, and $\succcurlyeq$ is intrinsically neutral to information π.

□

Proposition 1.3

Proof. By Theorem 1 in Strzalecki [2011], if $\succcurlyeq_0$ has a multiplier representation, then Savage's Sure-Thing principle is satisfied. So $\forall f \in F$ and x such that $fEx \sim_0 x$, we have $f \sim_0 xEf$. By step 1 of our proof for Theorem 1.2, this yields information neutrality. □

Theorem 1.3 and Corollary 1

Lemma A.4. *For the conditional cost function*

$$c_E(p_E) = \inf_{\{p\in\Delta(S):p(\cdot|E)=p_E\}} \frac{c(p)}{p(E)} \tag{A.1}$$

if $p(E) > 0$ for all $p \in c^{-1}(0)$, then the infimum attains at some $p \in \Delta(S)$, where $p(\cdot|E) = p_E$.

Proof. Let $Q(p_E) := \{p \in \Delta(S) : p(\cdot|E) = p_E\}$. Then $\overline{Q(p_E)} = Q(p_E) \cup \Delta(E^c)$ is compact in $\Delta(S)$. If $c(p) = +\infty$ for all $p \in Q(p_E)$, then $c_E(p_E) = +\infty$ and the infimum attains at any $p \in Q(p_E)$. Otherwise, $c_E(p_E) < +\infty$. By the definition of infimum, we can find a sequence $p^n \in Q(p_E)$, such that $\frac{c(p^n)}{p^n(E)}$ is decreasing and $\lim_n \frac{c(p^n)}{p^n(E)} = c_E(p_E)$. By compactness of $\overline{Q(p_E)}$, we can find a subsequence of $\{p^n\}$, say $\{p^k\}$, such that $p^k \to_k p^* \in \overline{Q(p_E)}$. It remains to show that if $p(E) > 0$ for all $p \in c^{-1}(0)$, then $p^* \notin \Delta(E^c)$.

Suppose not, so $p^* \in \Delta(E^c)$. By assumption, $c(p^*) > 0$, so $\frac{c(p^*)}{p^*(E)} = +\infty$. Yet by lower semicontinuity of c, $c_E(p_E) = \liminf_k \frac{c(p^k)}{p^k(E)} \geq \frac{c(p^*)}{p^*(E)} = +\infty$. A contradiction. □

From our discussion in Appendix A.2, Strong Monotonicity of $\succcurlyeq_0$ ensures that all events are $\succcurlyeq_0$-non-null. As a result, the condition that $p(E) > 0$ for all $p \in c^{-1}(0)$ is satisfied for all E.

We then verify that c_E is convex, lower semicontinuous and grounded, so c_E can serve as a cost function.

Lemma A.5. *The function $c_E : \Delta(E) \to [0, \infty]$ defined in (A.1) is (i) convex, (ii) lower semicontinuous, and (iii) grounded.*

Proof. Convexity. By the lower semicontinuity of c, $\forall p_E, q_E \in \Delta(E), \alpha \in [0,1]$, we can find $p^*, q^* \in \Delta$ such that $p^*(\cdot|E) = p_E, q^*(\cdot|E) = q_E$, and $c_E(p_E) = \frac{c(p^*)}{p^*(E)}, c_E(q_E) = \frac{c(q^*)}{q^*(E)}$. Fix $\alpha \in [0,1]$. Then there exists $\gamma \in [0,1]$ such that $\frac{\gamma p^*(E)}{\gamma p^*(E)+(1-\gamma)q^*(E)} = \alpha$. Set $p' := \gamma p^* + (1-\gamma)q^*$. Then $p'(\cdot|E) = \alpha p_E + (1-\alpha)q_E$. Therefore,

$$\begin{aligned} c_E(\alpha p_E + (1-\alpha)q_E) &\leq \frac{c(p')}{p'(E)} \leq \frac{\gamma c(p^*) + (1-\gamma)c(q^*)}{\gamma p^*(E) + (1-\gamma)q^*(E)} \\ &= \frac{\gamma p^*(E)}{\gamma p^*(E) + (1-\gamma)q^*(E)} c_E(p_E) + \frac{(1-\gamma)q^*(E)}{\gamma p^*(E) + (1-\gamma)q^*(E)} c_E(q_E) \\ &= \alpha c_E(p_E) + (1-\alpha)c_E(q_E). \end{aligned}$$

Lower semicontinuity. We want to show the epigraph $epi(c_E)$ is closed. To that end, let $(p_E^n, r_n) \in epi(c_E)$, $(p_E^n, r_n) \to_n (p_E, r)$. We want to show $r \geq c_E(p_E)$. Since $p(E) > 0$ for all $p \in c^{-1}(0)$, by the previous lemma $c_E(p_E^n) = \frac{c(p^n)}{p^n(E)}$ for some p^n where $p^n(\cdot|E) = p_E^n$. Since $\Delta(S)$ is compact, there exists a subsequence $\{p^k\}$ of $\{p^n\}$ such that $p_k \to_k p^*$.

If $p^*(E) > 0$, then $p^*(\cdot|E) = \lim_k p^k(\cdot|E) = \lim_k p_E^k = p_E$. Then $\liminf_k \frac{c(p^k)}{p^k(E)} \geq \frac{c(p^*)}{p^*(E)}$ by lower semicontinuity of c. Since $r_k \to r$ and $r_k \geq c_E(p_E^k) = \frac{c(p^k)}{p^k(E)}$, $r \geq \liminf_k \frac{c(p^k)}{p^k(E)} \geq \frac{c(p^*)}{p^*(E)} \geq c_E(p_E)$. Then we are done.

If $p^*(E) = 0$, then there must be a subsequence $p^k(E) \to_k 0$. Since $r + \epsilon \geq \tilde{c}_E(p_E^k) = \frac{c(p^k)}{p^k(E)}$ for $\epsilon > 0$ and sufficiently large k, $\liminf_k c(p^k) = 0 \geq c(p^*) \geq 0$. Thus $p^*(E) = 0$ and $c(p^*) = 0$, a contradiction.

Groundedness. c is grounded, so there exists p^* such that $c(p^*) = 0$. By assumption, $p^*(E) > 0$, so $c_E(p^*(\cdot|E)) = 0$. □

Lemma A.6. *Consider two variational functionals $I(\phi) = \min_{p \in \Delta} \langle \phi, p \rangle + c(p)$, and $I'(\phi) = \min_{p \in \Delta} \langle \phi, p \rangle + c'(p)$. If $c(p_0) < c'(p_0)$ for some p_0, then there exists $\xi \in B(\Sigma)$ such that $I(\xi) < I'(\xi)$.*

Proof. Consider the epigraph of c':

$$epi(c') = \{(p, r) \in \Delta \times \mathbb{R} | r \geq c'(p)\}$$

Since c' is nonnegative, convex, lower semicontinous, and grounded, $epi(c')$ is nonempty, closed and convex. Let $r_0 = c(p_0)$. Since $c(p_0) < c'(p_0)$, $(p_0, r_0) \notin epi(c')$. By the strict separating hyperplane theorem there exists $(\xi_0, r^*) \in B(\Sigma) \times \mathbb{R}$, $(\xi_0, r^*) \neq 0$, that strictly separates (p_0, r_0) from the set $epi(c')$, such that, that is

$$\langle \xi_0, p_0 \rangle + r_0 \cdot r^* < \inf_{r' \geq c'(p')} \langle \xi_0, p' \rangle + r' \cdot r^*$$

Note that we cannot have $r^* < 0$, otherwise we could take $r' = +\infty$ in the right hand side and the inequality fails. Also we cannot have $r^* = 0$, otherwise we get $\langle \xi_0, p_0 \rangle < \inf_{p'} \langle \xi_0, p' \rangle \leq \langle \xi_0, p_0 \rangle$, a contradiction. Thus $r^* > 0$, and we can rescale both sides by $\frac{1}{r^*}$ (take $\xi = \frac{1}{r^*}\xi_0$) to obtain

$$\langle \xi, p_0 \rangle + r_0 < \inf_{r' \geq c'(p')} \langle \xi, p' \rangle + r'$$

Then

$$\langle \xi, p_0 \rangle + r_0 = \langle \xi, p_0 \rangle + c(p_0) \geq \min_{p \in \Delta} \langle \xi, p \rangle + c(p) = I(\xi)$$

and

$$\inf_{r' \geq c'(p')} \langle \xi, p' \rangle + r' = \min_{p' \in \Delta} \langle \xi, p' \rangle + c'(p') = I'(\xi)$$

Thus $I(\xi) \leq \langle \xi, p_0 \rangle + r_0 < \inf_{r' \geq c'(p')} \langle \xi, p' \rangle + r' = I'(\xi)$. □

Proof of Theorem 1.3. (2) ⇒ (1). Suppose (2) holds. It is straightforward to verify Stable Risk Preferences and Consequentialism. We prove Conditional Certainty Equivalent Consistency also holds.

Fix $f \in \mathcal{F}$ and $x \in X$ such that $x \sim_E f$. We must prove $fEx \sim_0 x$. Suppose c and c_E satisfy update rule (A.1). Then

$$\begin{aligned} x \sim_E f \Rightarrow u(x) &= \inf_{p_E \in \Delta(E)} \int_E u(f) dp_E + c_E(p_E) \\ &= \inf_{p_E \in \Delta(E)} \int_E u(f) dp_E + \inf_{p \in \Delta : p(\cdot|E) = p_E} \frac{c(p)}{p(E)} \end{aligned}$$

Let $p^* \in \Delta$ achieve the infimum above.[1]

$$\begin{aligned} u(x) &= p^*(E) \left[\textstyle\int_E u(f) dp^*(\cdot|E) + \frac{c(p^*)}{p^*(E)} \right] + p^*(E^c) u(x) \\ &= \int_E u(f) dp^* + p^*(E^c) u(x) + c(p^*) \\ &\geq \min_{p \in \Delta} \int_E u(f) dp + p(E^c) u(x) + c(p) = V_0(fEx) \end{aligned}$$

[1] Let $I_E : B(\Sigma_E, u(X)) \to \mathbb{R}$ be such that $I_E(\xi) = \inf_{p_E \in \Delta(E)} \int_E \xi dp_E + c_E(p_E)$. Then I_E is also a variational functional. Applying Maccheroni et al. [2006a] Lemma 26, the infimum attains at some p_E^*. In addition, if $p(E) > 0$ for all $p \in c^{-1}(0)$, by the previous lemma there exists $p^* \in \Delta(S)$, $p^*(\cdot|E) = p_E^*$, at which the second infimum attains.

It remains to show that the inequality cannot be strict. If not, then $u(x) > V_0(fEx)$. Let $\tilde{p} \in argmin_{p\in\Delta} \int_E u(fEx)dp + p(E^c)u(x) + c(p)$. Then

$$\begin{aligned} u(x) > V_0(fEx) &= \min_{p\in\Delta} \int_E u(f)dp + p(E^c)u(x) + c(p) \\ &= \int_E u(f)d\tilde{p} + \tilde{p}(E^c)u(x) + c(\tilde{p}) \end{aligned}$$

If $\tilde{p}(E) = 0$, then $u(x) > u(x)+c(\tilde{p})$, which contradicts the non-negativity of c. So $\tilde{p}(E) > 0$. Then

$$\begin{aligned} u(x) &> \frac{1}{\tilde{p}(E)}[\int_E u(f)d\tilde{p} + c(\tilde{p})] \\ &= \int_E u(f)d\tilde{p}(\cdot|E) + \frac{c(\tilde{p})}{\tilde{p}(E)} \\ &\geq \min_{p_E\in\Delta(E)} \int_E u(f)dp_E + \min_{p\in\Delta:p(\cdot|E)=p_E} \frac{c(p)}{p(E)} \\ &= V_E(f) \end{aligned}$$

This contradicts the assumption that $x \sim_E f$. So $fEx \sim_0 x$.

For the converse, suppose $fEx \sim_0 x$. Then

$$\begin{aligned} u(x) = V_0(fEx) &= \min_{p\in\Delta(S)} \int_E u(f)dp + u(x)p(E^c) + c(p) \\ &= \int_E u(f)dp^* + u(x)p^*(E^c) + c(p^*) \end{aligned}$$

where $p^* \in argmin_p \int_E u(f)dp + u(x)p(E^c) + c(p)$. If $p^*(E) = 0$, then the equality above implies $c(p^*) = 0$, a contradiction to the assumption that $p(E) > 0, \forall p \in c^{-1}(0)$. So $p^*(E) > 0$, and

$$p^*(E)u(x) = \int_E u(f)dp^* + c(p^*)$$

Thus

$$\begin{aligned} u(x) &= \int_E u(f)dp^*(\cdot|E) + \frac{c(p^*)}{p^*(E)} \\ &\geq \min_{p_E} \int_E u(f)dp_E + \inf_{p\in\Delta:p(\cdot|E)=p_E} \frac{c(p)}{p(E)} = V_E(f) \end{aligned}$$

So $x \succcurlyeq_E f$.

Also, as argued before, we can find $q^* \in \Delta(S)$, $q^*(E) > 0$, such that $V_E(f) = \int_E u(f)dq^*(\cdot|E) + \frac{c(q^*)}{q^*(E)}$. So

$$q^*(E)\left[\int_E u(f)dq^*(\cdot|E) + \tfrac{c(q^*)}{q^*(E)}\right] + q^*(E^c)u(x) \geq V_0(fEx) = u(x)$$

Thus $V_E(f) \geq u(x)$, or $f \succcurlyeq_E x$. So $x \sim_E f$.

(1) $\Rightarrow$ (2). By assumption, $\succcurlyeq_E$ has a representation of the form

$$V_E(f) = \min_{p \in \Delta(S)} \int_S u_E(f)dp + c_E(p)$$

By Stable Risk Preferences, $\succcurlyeq_0$ and $\succcurlyeq_E$ agree on constant acts X. We can normalize by setting $u_E = u$. Next we want to show only p with support on E can achieve the minimum defining V_E. For each $f \in \mathcal{F}$, choose $p^* \in \arg\min_{p \in \Delta(S)} \int_S u(f)dp + c_E(p)$. Without loss of generality, we can choose $x_* \in X$ such that $f(s) \succ_0 x_*$ for all s.[2] Since $(fEx_*)Ex = fEx$ for any x, by Conditional Certainty Equivalent Consistency, $fEx_* \sim_E f$. Then

$$V_E(f) = \int_S u(f)dp^* + c_E(p^*) = V_E(fEx_*) \leq \int_E u(f)dp^* + p^*(E^c)u(x_*) + c_E(p^*)$$

So $\int_{E^c}(u(f) - u(x_*))dp^* \leq 0$. Since $u(f) - u(x_*)$ is strictly positive on E^c, $\int_{E^c}(u(f) - u(x_*))dp^* \geq 0$, and this is an equality if and only if $p^*(E^c) = 0$. So $p^*(E) = 1$, and p^* has a natural imbedding in $\Delta(E)$. Therefore $\forall f$,

$$V_E(f) = \min_{p \in \Delta(E)} \int_E u(f)dp + c_E(p)$$

It remains to show that the (unique) conditional cost function c_E coincides with $\tilde{c}_E(p_E) := \inf_{p \in \Delta: p(\cdot|E) = p_E} \frac{c(p)}{p(E)}$. Suppose not, so $c_E \neq \tilde{c}_E$. Thus there exists p_E^* such that $c_E(p_E^*) \neq \tilde{c}_E(p_E^*)$. We prove a contradiction for the case $c_E(p_E^*) > \tilde{c}_E(p_E^*)$. The case $c_E(p_E^*) < \tilde{c}_E(p_E^*)$ can be proved by replicating the arguments. Applying Lemma A.6, we can find $\xi_E \in B(\Sigma_E)$ such that $\min_{p_E} \int_E \xi_E dp_E + \tilde{c}_E(p_E) < \min_{p_E} \int_E \xi_E dp_E + c_E(p_E)$. Since $u(X)$ is unbounded, $B(\Sigma_E) \subseteq B(\Sigma_E, u(X)) + \mathbb{R}$. Thus there is an act $f \in \mathcal{F}$ such that $(u(f) + k)(s) = \xi_E(s)$ on E for some constant k. So $\min_{p_E} \int_E u(f)dp_E + \tilde{c}_E(p_E) < \min_{p_E} \int_E u(f)dp_E + c_E(p_E)$.

By Continuity, we can find $x \in X$ that is the E-conditional equivalent of f, $x \sim_E f$, and $u(x) = V_E(f) = \min_{p_E} \int_E u(f)dp_E + c_E(p_E)$.

Then

$$\begin{aligned}
u(x) &= \min_{p_E} \int_E u(f)dp_E + c_E(p_E) \\
&> \min_{p_E} \int_E u(f)dp_E + \tilde{c}_E(p_E) \\
&= \min_{p_E} \int_E u(f)dp_E + \inf_{p \in \Delta: p(\cdot|E) = p_E} \frac{c(p)}{p(E)} \\
&= \min_{p_E} \inf_{p \in \Delta: p(\cdot|E) = p_E} \int_E u(f)dp_E + \frac{c(p)}{p(E)} \\
&= \inf_{p \in \Delta, p(E) > 0} \frac{1}{p(E)} \left[\textstyle\int_E u(f)dp + c(p)\right]
\end{aligned}$$

[2] If not, then $u(X)$ is bounded below and $\min_s u(f)(s)$ achieves the lower bound. By translation invariance p^* is also a minimizing probability for f' such that $u(f') = u(f) + \epsilon$. Then the whole argument works for f'.

As argued before, we can find $\underline{p} \in argmin_{p\in\Delta,p(E)>0}\frac{1}{p(E)}\left[\int_E u(f)dp + c(p)\right]$. Then multiplying both sides of the inequality by $\underline{p}(E)$ and adding $\underline{p}(E^c)u(x)$ to both sides yields

$$\begin{aligned} u(x) &> \underline{p}(E)\left(\tfrac{1}{\underline{p}(E)}[\textstyle\int_E u(f)d\underline{p} + c(\underline{p})]\right) + \underline{p}(E^c)u(x) \\ &= \int_E u(f)d\underline{p} + \underline{p}(E^c)u(x) + c(\underline{p}) \\ &= \int u(fEx)d\underline{p} + c(\underline{p}) > V_0(fEx) \end{aligned}$$

So $x \succ_0 fEx$, violating Conditional Certainty Equivalent Consistency. ☐

Proof of Corollary 1. For part (1), suppose $\succcurlyeq_0$ has a MEU representation $(u, \mathcal{P})$. So it has a variational representation (u, c) with cost function c such that $c(p) = 0$ if $p \in \mathcal{P}$ and $c(p) = +\infty$ if $p \notin \mathcal{P}$. For any nonempty event E, Strong Monotonicity of $\succcurlyeq_0$ ensures that $p(E) > 0$ for all $p \in \mathcal{P}$. Applying updating rule A.1,

$$c_E(p_E) = \begin{cases} 0 \text{ if } p_E \in \mathcal{P}_E = \{p(\cdot|E)|p \in \mathcal{P}\} \\ +\infty \text{ otherwise} \end{cases}$$

So $\succcurlyeq_E$ has MEU representation $(u, \mathcal{P}_E)$.

For part (2), suppose $\succcurlyeq_0$ also has a multiplier preference representation (u, q, θ). So it has a variational representation (u, c) with cost function $c(p) = \theta \int ln\frac{p}{q}dp$. For any nonempty event E, Strong Monotonicity of $\succcurlyeq_0$ ensures that $q(E) > 0$. Applying updating rule A.1,

$$\begin{aligned} c_E(p_E) &= \min_{p\in\Delta(S):p(\cdot|E)=p_E} \frac{\theta}{p(E)}\int ln\frac{p}{q}dp \\ &= \min_{p\in\Delta(S):p(\cdot|E)=p_E} \frac{\theta}{p(E)}[(\int_E ln\frac{p_E}{q_E}dp_E)p(E) + (\int_{E^c} ln\frac{p_{E^c}}{q_{E^c}}dp_{E^c})p(E^c) \\ &\qquad + (p(E)ln\frac{p(E)}{q(E)} + p(E^c)\ln\frac{p(E^c)}{q(E^c)})] \\ &= \theta\int_E ln\frac{p_E}{q_E}dp_E \end{aligned}$$

In the last step, we choose p such that $p(E) = q(E)$ and $p(\cdot|E^c) = q(\cdot|E^c)$. So $\succcurlyeq_E$ has multiplier representation (u, q_E, θ). ☐

Proposition 1.4

Proof. By Theorem 1.2, $\succcurlyeq$ exhibits intrinsic information aversion at all acts if and only if $\forall f$ and x such that $fEx \sim_0 x$, $f \succcurlyeq_0 xEf$. By Conditional Certainty Equivalent Consistency, $fEx \sim_0$ if and only if $x \sim_E f$.

If $x \sim_E f$, then $u(x) = \min_{p_E \in \Delta(E)} \int_E u(f)dp_E + c_E(p_E)$. So

$$\begin{aligned}
V_0(xEf) &= \min_{p\in\Delta} p(E)u(x) + \int_{E^c} u(f)dp + c(p) \\
&= \min_{p\in\Delta} p(E)[\min_{p_E\in\Delta(E)} \int_E u(f)dp_E + \hat{c}_E(p_E)] + \int_{E^c} u(f)dp + c(p) \\
&= \min_{p\in\Delta} \min_{p_E\in\Delta(E)} p(E)[\int_E u(f)dp_E + \hat{c}_E(p_E)] + + \int_{E^c} u(f)dp + c(p) \\
&= \min_{q\in\Delta} \min_{q_E\in\Delta(E)} \int u(f)dq + q(E)\hat{c}_E(q_E) + c(q_E \otimes_E q) \\
&\text{(change of variable: } q = p_E \otimes_E p \text{, and } q_E = p(\cdot|E)) \\
&= \min_{q\in\Delta} \int u(f)dq + q(E)\hat{c}_E(q_E) + \min_{q_E\in\Delta(E)} c(q_E \otimes_E q)
\end{aligned}$$

Also

$$V_0(f) = \min_{q\in\Delta} \int u(f)dq + c(q)$$

"If" direction. Suppose $\inf_{q_E\in\Delta(E)} c(q_E \otimes_E p) + p(E) \inf_{q\in\Delta(S)} \frac{c(p_E\otimes q)}{q(E)} \leq c(p), \forall p$. Then for all f, q,

$$\int u(a)dq + q(E)\hat{c}_E(q_E) + \min_{q_E\in\Delta(E)} c(q_E \otimes_E q) \leq \int u(f)dq + c(q),$$

so $V_0(xEf) \leq V_0(f)$. Thus the DM is averse to partial information at all f.

"Only if" direction. For each $E \in \Sigma$, define

$$\tilde{c}(p) = \begin{cases} \inf_{q_E\in\Delta(E)} c(q_E \otimes_E p) + p(E)c_E(p(\cdot|E)) & \text{if } p(E) > 0 \\ +\infty & \text{otherwise} \end{cases}$$

Define $\tilde{I} : B(S, \mathbb{R}) \to \mathbb{R}$ by $\tilde{I}(\xi) = \inf_{p\in\Delta(S)} \int \xi dp + \tilde{c}(p)$. By the calculation above, we have $\forall f \in \mathcal{F}$, $x \sim_E f$, $V_0(xEf) = \tilde{I}(u(f))$.

If statement (2) fails, then there exists p such that $\tilde{c}(p) > c(p)$. By Lemma A.6, we can find $\xi \in B(S, \mathbb{R})$ such that $\tilde{I}(\xi) > I(\xi)$. By unboundedness, $B(S, \mathbb{R}) \subseteq B(S, u(X)) + \mathbb{R}$, so there exists $f \in F$ such that $u(f)+k = \xi$ for some constant k. So we can find $f \in \mathcal{F}$ such that $V_0(xEf) = \tilde{I}(u(f)) > I(u(f)) = V_0(f)$. This contradicts aversion to partial information. $\square$

Proposition 1.5

Proof. Let $p^* \in c^{-1}(0) \cap \arg\min_{p\in\Delta}[\int_S u(f)dp + c(p)]$. Then $\forall \pi = \{E_1, \cdots, E_n\}$,

$$\begin{aligned} V(\pi, f) &= \min_{p\in\Delta} \sum p(E_i) \left[\min_{p_i\in\Delta(E_i)} \int u(f)dp_i + c_{E_i}(p_i)\right] + \min_{\{q\in\Delta(S):q=p \text{ on } \pi\}} c(q) \\ &\leq \sum p^*(E_i) \left[\int u(f)dp^*(\cdot|E_i) + c_{E_i}(p^*(\cdot|E_i))\right] + \min_{\{q\in\Delta(S):q=p^* \text{ on } \pi\}} c(q^*) \\ &= \int_S u(f)dp^* \\ &= \int_S u(f)dp^* + c(p^*) = V(\pi_0, f) \end{aligned}$$

The second equality follows from

$$c_{E_i}(p^*(\cdot|E_i)) = \min_{p(\cdot|E_i)=p^*(\cdot|E_i)} \frac{c(p)}{p(E_i)} = 0$$

and

$$\min_{\{q\in\Delta(S):q=p^* \text{ on } \pi\}} c(q^*) = 0.$$

□

Proposition 1.8

Proof. (1) Suppose $\succcurlyeq_0$ has a CEU representation (u, ν) and satisfies Uncertainty Aversion. By the Proposition in Schmeidler [1989] the corresponding functional I_0 is concave and superadditive. By Proposition 1.1, this implies that Event Complementarity holds. By Theorem 1.2, $\succcurlyeq$ exhibits aversion to partial information.

(2) Suppose $\succcurlyeq_0$ has a CEU representation (u, ν) and satisfies Uncertainty Loving. By Schmeidler [1989] Remark 6 the corresponding functional I_0 is convex and subadditive. By Proposition 1.1 and Theorem 1.2, $\succcurlyeq$ exhibits attraction to partial information. □

Proposition 1.9

Proof. Fix $\pi = \{E_1, \cdots, E_n\}$. Suppose $\succcurlyeq_0$ has second order belief representation $(u, \phi; \Theta, \mu)$ and $\succcurlyeq_0$ is ambiguity averse. Then by Klibanoff et al. [2005] Proposition 1, ϕ is concave. Let

f be an act where $\succcurlyeq_0$ displays local ambiguity neutrality. Then

$$\begin{aligned}
V(\pi, f) &= \int_\Theta \phi[\sum_{i=1}^n p_{\theta'}(E_i)\phi^{-1}[\int_\Theta \phi(\int u(f)dp_{\theta_i}(\cdot|E_i))d\mu_{E_i}(\theta_i)]]d\mu(\theta') \\
&\leq \int_\Theta \phi[\sum_{i=1}^n p_{\theta'}(E_i)[\int_\Theta \int u(f)dp_{\theta_i}(\cdot|E_i)d\mu_{E_i}(\theta_i)]]d\mu(\theta') \\
&= \int_\Theta \phi[\sum_{i=1}^n p_{\theta'}(E_i)(\int_\Theta \int_{E_i} u(f)dp_{\theta_i}\frac{d\mu(\theta_i)}{\int p_{\theta''}(E_i)d\mu(\theta'')}]d\mu(\theta') \\
&= \int_\Theta \phi[\sum_{i=1}^n (\int_\Theta \int_{E_i} u(f)dp_{\theta_i}d\mu(\theta_i)\frac{p_{\theta'}(E_i)}{\int p_{\theta''}(E_i)d\mu(\theta'')}d\mu(\theta')] \\
&\leq \phi \int_\Theta [\sum_{i=1}^n (\int_\Theta \int_{E_i} u(f)dp_{\theta_i}d\mu(\theta_i)\frac{p_{\theta'}(E_i)}{\int p_{\theta''}(E_i)d\mu(\theta'')}d\mu(\theta')] \\
&= \phi[(\sum_{i=1}^n \int_\Theta \int_{E_i} u(f)dp_{\theta_i}d\mu(\theta_i)(\int_\Theta \frac{p_{\theta'}(E_i)}{\int p_{\theta''}(E_i)d\mu(\theta'')}d\mu(\theta'))] \\
&= \phi(\int_\Theta \int_S u(f)dp_\theta d\mu) = V(\pi_0, f)
\end{aligned}$$

The two inequalities follow from the concavity of ϕ. The last equality holds because $\succcurlyeq_0$ displays local ambiguity neutrality at f.

The case for ambiguity loving $\succcurlyeq_0$ can be proved analogously. □

Appendix B

Proofs for Chapter 2

Lemma 2.1

Proof. We first show $\succcurlyeq^+$ is represented by $\tilde{V}(\pi, F) = \max_{f \in F^\pi} V(\pi, f)$. For all (π, F) and (π', G), $(\pi, F) \succcurlyeq^+ (\pi', G)$ if and only if

$$\forall g \in G^{\pi'}, \exists f \in F^\pi, (\pi, f) \succcurlyeq (\pi', g)$$

Since $V : \Pi \times \mathcal{F}$ represents $\succcurlyeq$, this is equivalent to

$$\max_{f \in F^\pi} V(\pi, f) \geq \max_{g \in G^{\pi'}} V(\pi', g)$$

Thus $(\pi, F) \succcurlyeq^+ (\pi', G)$ if and only if $\tilde{V}(\pi, F) \geq \tilde{V}(\pi', G)$.

Then we show $\max_{f \in F^\pi} V(\pi, f) = V_0(c(F|\pi))$. By definition, $F^\pi = \{f_1 E_1 f_2 E_2 \cdots E_{n-1} f_n : f_i \in F, \forall i = 1, \cdots, n\}$. So

$$\begin{aligned}
\max_{f \in F^\pi} V(\pi, f) &= \max_{f_1 \in F} \ldots \max_{f_n \in F} V(\pi, f_1 E_1 \cdots E_{n-1} f_n) \\
&= \max_{f_1 \in F} \ldots \max_{f_n \in F} V(\pi, [c(f_i|E_i), E_i]_1^n) \quad \text{by } \pi\text{-Recursivity} \\
&= \max_{f_1 \in F} \ldots \max_{f_n \in F} V_0([c(f_i|E_i), E_i]_1^n) \\
&= V_0([c(F|E_i), E_i]_1^n) \quad \text{by } \pi_0\text{-monotonicity}
\end{aligned}$$

where the second to last equality is due to Independence from Redundant Information and Time Neutrality. □

Proposition 2.2 and 2.3

Proof of Proposition 2.2. (2) ⇔ (3) is due to Theorem 1.2.

To show (1) ⇒ (2), take any singleton menu $F = \{f\}$. A preference for perfect information implies $(\pi^*, f) \succcurlyeq (\pi, f), \forall \pi$. By Time Neutrality, $(\pi^*, f) \sim (\pi_0, f)$, so $(\pi_0, f) \succcurlyeq (\pi, f)$.

To show (2) $\Rightarrow$ (1). Let $\pi \in \Pi$ and $F \in \mathcal{M}$. Then

$$V(\pi^*, F) - V(\pi, F) = [\max_{f\in F^{\pi^*}} V(\pi^*, f) - \max_{f\in F^{\pi}} V(\pi^*, f)] + [\max_{f\in F^{\pi}} V(\pi^*, f) - \max_{f\in F^{\pi}} V(\pi, f)]$$

The first term is non-negative since $F^\pi \subseteq F^{\pi^*}$. By (2) and Time Neutrality, $V(\pi^*, f) = V(\pi_0, f) \geq V(\pi, f)$, for all π, f. So

$$max_{f\in F^\pi} V(\pi, f) = V(\pi, f^*) \leq V(\pi^*, f^*) \leq \max_{f\in F^\pi} V(\pi^*, f)$$

where $f^* \in F^\pi$ is the act that maximizes $V(\pi, \cdot)$. So the second term is also non-negative. Thus $V(\pi^*, F) \geq V(\pi, F)$ and the DM has preferences for perfect information. □

Next we prove Proposition 2.3. We first prove a lemma. Let $F_0 = \arg\max_{f\in F} V(\pi_0, f)$ be the set of uninformed optimal acts. By our decomposition, as long as the DM is not strictly averse to information π at some $f_0 \in F_0$, then information is valuable.

Let $F_i^* = \arg\max_{f\in F} V_{E_i}(f)$ be the set of optimal acts in F conditional on learning about E_i. Consider $F^* = \{f_1^* E_1 f_2^* E_2 \cdots E_{n-1} f_n^* : f_i^* \in F_i^*, \forall i\} \subseteq F^\pi$. The instrumental value of information is zero if and only if $F^* \cap F \neq \emptyset$. We collect these observations below.

Lemma B.1. *1. If there exists an unconditional optimal act $f_0 \in F_0$ such that $V(\pi, f_0) \geq V(\pi_0, f_0)$ at f_0, then $V(\pi, F) - V(\pi_0, F) \geq 0$.*

2. If there exists a conditional optimal strategy $f^* \in F^*$ such that $f^* \in F$ and $V(\pi, f^*) \leq (<) V(\pi_0, f^*)$, then $V(\pi, F) - V(\pi_0, F) \leq (<) 0$.

Proof. By definition $V(\pi_0, f_0) = \max_{f\in F} V(\pi_0, f)$. If $V(\pi, f_0) \geq V(\pi_0, f_0)$, then the intrinsic value of information π at menu F is non-negative:

$$\max_{f\in F} V(\pi, f) - \max_{f\in F} V(\pi_0, f) \geq V(\pi, f_0) - V(\pi_0, f_0) \geq 0.$$

As the instrumental value of information is always non-negative, $V(\pi, F) - V(\pi_0, F) \geq 0$ and π is valuable.

If there exists $f^* \in F \cap F^*$, then the instrumental value of π, $V(\pi, f^*) - \max_{f\in F} V(\pi, f) = 0$. In addition $\max_{f\in F} V(\pi, f) = V(\pi, f^*) \leq V(\pi_0, f^*) \leq \max_{f\in F} V(\pi_0, f)$, so the intrinsic value of π is non-positive. □

Remark 5. The first condition is helpful, as it requires only calculation of an optimal act in the uninformed case. This could simplify checking whether ambiguity aversion generates information aversion or not. In MEU models, this is equivalent to $V(\pi, f_0) = V(\pi_0, f_0)$, when the intrinsic value of information π for menu F vanishes.

Proof of Proposition 2.3. If there exists an uninformed optimal act f_0 that is π-measurable, then $V(\pi, f_0) = V(\pi^*, f_0) = V(\pi_0, f_0)$. By the above lemma, $\Delta V(\pi, F) \geq 0$. □

Proof of Corollary 2. Let x be the uninformed optimal act for DM 1. So $V^1(\pi_0, x) \geq V^1(\pi_0, f)$, for all f in menu F. Since DM 2 is more ambiguity averse than DM 1, $u_2 = u_1$ and $c_2 \leq c_1$. So for all $f \in \mathcal{F}$,

$$V^2(\pi_0, f) = \min_{p \in \Delta(S)} \int_S u(f)dp + c_2(p) \leq \min_{p \in \Delta(S)} \int_S u(f)dp + c_1(p) = V^1(\pi_0, f)$$

and $V^1(\pi_0, x) = u(x) = V^2(\pi_0, x)$. Thus $V^2(\pi_0, x) \geq V^2(\pi_0, f)$ for all $f \in F$. Since x is π-measurable, by Proposition 2.3 we have $\Delta V^2(\pi, F) \geq 0$. □

Marginal Value of Information

For any menu F, consider two partitions $\pi_2 \geq \pi_1$. The marginal value of getting the finer information π_2 is:

$$V(\pi_2, F) - V(\pi_1, F) = [\max_{f \in F^{\pi_2}} V(\pi_2, f) - \max_{f \in F^{\pi_1}} V(\pi_2, f)] + [\max_{f \in F^{\pi_1}} V(\pi_2, f) - \max_{f \in F^{\pi_1}} V(\pi_1, f)]$$

The first term captures the instrumental value of getting finer information π_2 relative to π_1, and since $F^{\pi_1} \subseteq F^{\pi_2}$ this term is non-negative. The second part captures the intrinsic value of information π_2 relative to π_1.

Lemma B.1 can be generalized as follows.

Lemma B.2. *1. If there exists an optimal strategy f^{*1} for decision problem (π_1, F) such that $V(\pi_1, f^{*1}) \leq V(\pi_2, f^{*1})$, then $V(\pi_2, F) - V(\pi_1, F) \geq 0$.*

2. If there exists an optimal strategy f^{*2} for decision problem (π_2, F) such that $f^{*2} \in F^{\pi_1}$ and $V(\pi_1, f^{*2}) \geq V(\pi_2, f^{*2})$, then $V(\pi_2, F) - V(\pi_1, F) \leq 0$.

The proof is similar to the proof of Lemma B.1 and hence omitted.

Appendix C

Proofs for Chapter 3

Properties of the Multiple Priors Utility

Lemma C.1. *For all (C, ∞)-bandit, strategy $\mathbf{a}$, and history h_t, $\{U_{h_t}^T(C, \mathbf{a})\}_{T=t+1}^{\infty}$ converges uniformly to some $U_{h_t}(C, \mathbf{a})$.*

Proof. For any fixed h_t, C and $\mathbf{a}$, let $Z = (Z_1, \dots, Z_t, \cdots)$ be the resulting infinite random payoff stream, and $Z^T = (Z_1, \dots, Z_T)$ be corresponding T-truncation. First we show for every $C, \mathbf{a}$ and h_t, $\{U_{h_t}^T(C, \mathbf{a})\}_{T=t+1}^{\infty}$ is a Cauchy sequence in $\mathbb{R}$ and thus converges.

For any $T_1 > T_2 > T$, think of a large T, let

$$\Delta := U_{h_t}^{T_1}(C, \mathbf{a}) - U_{h_t}^{T_2}(C, \mathbf{a}) = U_{h_t}^{T_1}(Z_1, \dots, Z_{T_2}, Z_{T_2+1}, \dots, Z_{T_1}) - U_{h_t}^{T_1}(Z_1, \dots, Z_{T_2}, 0, \dots, 0)$$

Since $|Z_t| \leq M$ and $U_{h_t}^{T_1}$ is monotone in the payoff streams,

$$U_{h_t}^{T_1}(0, \dots, 0, -M, \dots, -M) \leq \Delta \leq U_{h_t}^{T_1}(0, \dots, 0, M, \dots, M)$$

where in both sides the first T_2 coordinates are zero. Thus $|\Delta| < \delta^T \frac{M}{1-\delta}$. For arbitrary $\epsilon > 0$, we can pick T large enough that $|\Delta| < \epsilon$. So the sequence $\{U_{h_t}^T(C, \mathbf{a})\}$ is Cauchy and converges. Let the limit be $U_{h_t}(C, \mathbf{a})$. Finally since for arbitrary ϵ, the selection of T does not depend on C or $\mathbf{a}$, so the sequence of functions $\{U_{h_t}^T\}_{T=t+1}^{\infty}$ converges uniformly to function U_{h_t}. □

Lemma C.2. *For all one-armed bandit (C, λ, T), the value function $V(C, \lambda, T)$ is continuous, non-decreasing, and convex in λ.*

Proof. By Theorem 3.2.3, the value function satisfies

$$V(C, \lambda, T) = \frac{1 - \delta^T}{1 - \delta} \max\{\Lambda(C, T), \lambda\}$$

Since $\Lambda(C, T)$ does not depend on λ, the lemma follows as a direct consequence. □

Proofs

Proof of Fact 1. For finite T, we prove by induction on T. For $T = 1$,

$$V(\mu_{\tau,\sigma^2}, \lambda, 1) = E_{\mu_{\tau,\sigma^2}}[X_1^1] \vee \lambda = \tau \vee \lambda$$

V is weakly increasing in τ and σ^2, and is a convex function of τ.

Now suppose the claim is true for all horizon less than T. We want to show it is true for T as well. To see this, note that

$$V(\mu_{\tau,\sigma^2}, \lambda, T) = \max\{\lambda\frac{1-\delta^T}{1-\delta}, E_{\mu_{\tau,\sigma^2}}[X_1^1 + V(\mu_{\frac{\tau+X_1^1}{2},\sigma^2}, \lambda, T-1)]\}$$

By induction hypothesis, $V(\mu_{\frac{\tau+X_1^1}{2},\sigma^2}, \lambda, T-1)$ is increasing in $\frac{\tau+X_1^1}{2}$ and σ^2, and is a convex function of $\frac{\tau+X_1^1}{2}$. Thus $X_1^1 + V(\mu_{\frac{\tau+X_1^1}{2},\sigma^2}, \lambda, T-1)$ is increasing in τ, X_1^1, and σ^2, and is a convex function of X_1^1 and τ. Since μ_{τ,σ^2} is increasing in μ in first order stochastic dominance ranking, and increasing in σ^2 in second order stochastic dominance ranking, $V(\mu_{\tau,\sigma^2}, \lambda, T)$ is weakly increasing in τ and σ^2, and is a convex function of τ.

For infinite horizon problems, note that the value function can be approximated by the corresponding finite horizon value function, so monotonicity and convexity is preserved in the limit. □

Proof of equation (3.6) and (3.7). We prove (3.6) by induction on T. For $T = 0$, (3.6) is true vacuously. Suppose it is also true for all problems with horizon less than T, we want to show that it is true for problems with horizon T. So

$$V(C_{a,b,\sigma^2}, \lambda, T) = \max\{\lambda\frac{1-\delta^T}{1-\delta}, \inf_{\tau\in[a,b]} E_{\mu_{\tau,\sigma^2}}[X_1^1 + \delta V(C_{a,b,\sigma^2}(\cdot|X_1^1), \lambda, T-1)]\}$$

, where

$$\begin{aligned} V(C_{a,b,\sigma^2}(\cdot|X_1^1), \lambda, T-1) &= V(C_{\frac{a+X_1^1}{2},\frac{b+X_1^1}{2},\frac{\sigma^2}{2}}, \lambda, T-1) \\ &= V(\mu_{\frac{a+X_1^1}{2},\frac{\sigma^2}{2}}, \lambda, T-1) \end{aligned}$$

by induction hypothesis. Also since $X_1^1 + \delta V(\mu_{\frac{a+X_1^1}{2},\frac{\sigma^2}{2}}, \lambda, T-1)$ is increasing in X_1^1,

$$\begin{aligned} V(C_{a,b,\sigma^2}, \lambda, T) &= \max\{\lambda\frac{1-\delta^T}{1-\delta}, E_{\mu_{a,\sigma^2}}[X_1^1 + \delta V(\mu_{\frac{a+X_1^1}{2},\frac{\sigma^2}{2}}, \lambda, T-1)]\} \\ &= V(\mu_{a,\sigma^2}, \lambda, T) \end{aligned}$$

When $T = \infty$, both value functions can be approximated by corresponding finite horizon value functions, and equality still holds.

For (3.7), $\Lambda(C_{a,b,\sigma^2}, T)$ is the unique cutoff λ that solves

$$\lambda\frac{1-\delta^T}{1-\delta} = \inf_{\tau\in[a,b]} E_{\mu_{\tau,\sigma^2}}[X_1^1 + \delta V(C_{a,b,\sigma^2}(\cdot|X_1^1), \lambda, T-1)]$$

By (3.6), this is equivalent to the λ solving

$$\lambda\frac{1-\delta^T}{1-\delta} = E_{\mu_{a,\sigma^2}}[X_1^1 + \delta V(\mu_{a,\sigma^2}(\cdot|X_1^1), \lambda, T-1)]$$

and by definition the latter solution is $\Lambda(\mu_{a,\sigma^2}, T)$. So (3.7) holds.

□

Proof of Corollary 5. For (1) ⇒ (2), note that by Lemma 2 in Li [2012], $\max_N U(C, \lambda, N, T) = V(C, \lambda, T)$ and $\min_{\mu\in C} \max_N U(\mu, \lambda, N, T) = \min_{\mu\in C} V(\mu, \lambda, T)$. By Theorem 3.2.3,

$$V(C, \lambda, T) = \frac{1-\delta^T}{1-\delta}\max\{\Lambda(C, T), \lambda\}$$

$$\min_{\mu\in C} V(\mu, \lambda, T) = \frac{1-\delta^T}{1-\delta}\min_{\mu\in C}\max\{\Lambda(\mu, T), \lambda\} = \frac{1-\delta^T}{1-\delta}\max\{\min_{\mu\in C}\Lambda(\mu, T), \lambda\}$$

If (1) also holds, then

$$\max_N U(C, \lambda, N, T) = \min_{\mu\in C}\max_N U(\mu, \lambda, N, T) = \frac{1-\delta^T}{1-\delta}\max\{\Lambda(C, T), \lambda\}$$

For (2) ⇒ (1), let $\lambda = \Lambda(C, T)$. Then by definition $V(C, \lambda, T) = \frac{1-\delta^T}{1-\delta}\Lambda(C, T)$. Also

$$\begin{aligned}\min_{\mu\in C} V(\mu, \lambda, T) &= \frac{1-\delta^T}{1-\delta}\min_{\mu\in C}\max\{\Lambda(\mu, T), \Lambda(C, T)\}\\ &= \frac{1-\delta^T}{1-\delta}\min_{\mu\in C}\Lambda(\mu, T)\end{aligned}$$

where the second equality is by Corollary 4.

If (2) holds, then $V(C, \lambda', T) = \min_{\mu\in C} V(\mu, \lambda', T)$ for all λ'. As a result, $\Lambda(C, T) = \min_{\mu\in C}\Lambda(\mu, T)$. □

Proof of Proposition 3.5. Suppose C and Θ in one-armed bandit (C, λ, T) satisfy Condition 1. We first show that if C and Θ satisfy Condition 1, then for any observation of $s \in S$, $\{\mu_a(\cdot|s) : a \in [\underline{a}, \bar{a}]\}$ also has MLRP in θ. To see this, note that

$$\mu_a(\theta|s) = \frac{\mu_a(\theta)l(s|\theta)}{\int_\Theta l(s|\theta')d\mu_a(\theta')}$$

for all a. So for $a_1 < a_2, \theta_1 < \theta_2$, substituting the above equality

$$\frac{\mu_{a_2}(\theta_2|s)}{\mu_{a_2}(\theta_1|s)} = \frac{\mu_{a_2}(\theta_2)l(s|\theta_2)}{\mu_{a_2}(\theta_1)l(s|\theta_1)} \geq \frac{\mu_{a_1}(\theta_2)l(s|\theta_2)}{\mu_{a_1}(\theta_1)l(s|\theta_1)} = \frac{\mu_{a_1}(\theta_2|s)}{\mu_{a_1}(\theta_1|s)}$$

where the inequality is due to $\{\mu_a \in \Delta(\Theta) : a \in [\underline{a}, \bar{a}]\}$ has MLRP in θ. Therefore the MLRP is preserved after Bayesian updating.

Let $P_{\mu_a}(\cdot) = \int_\Theta l(\cdot|\theta)d\mu_a(\theta)$ be the predictive distribution on X induced by Bayesian prior μ_a. We show that for all $\underline{a} \leq a_1 < a_2 \leq \bar{a}$, $P_{\mu_{a_2}}$ FOSD $P_{\mu_{a_1}}$. For any $x \in \mathbb{R}$,

$$P_{\mu_{a_1}}(X_1 \leq x) = \int_\Theta P(X_1 \leq x|\theta)d\mu_{a_1}(\theta) \geq \int_\Theta P(X_1 \leq x|\theta)d\mu_{a_2}(\theta) = P_{\mu_{a_2}}(X_1 \leq x)$$

where the inequality is due to that $P(X_1 \leq x|\theta)$ is weakly decreasing in θ and μ_{a_2} FOSD μ_{a_1}.

Fix $a \in [\underline{a}, \bar{a}]$ and $x_1 < x_1'$, for arbitrary subsequent history $(x_2, \dots, x_t)$ $(1 \leq t \leq T)$, we have

$$\frac{\mu_a(\theta|x_1', x_2, \dots, x_t)}{\mu_a(\theta|x_1, x_2, \dots, x_t)} = \frac{l(x_1'|\theta)}{l(x_1|\theta)} \cdot \frac{\mu_a(\theta)l(x_2, \dots, x_t|\theta)}{\mu_a(\theta)l(x_2, \dots, x_t|\theta)} \cdot \frac{\int_\Theta l(x_1, x_2, \dots, x_t|\theta')d\mu_a(\theta')}{\int_\Theta l(x_1', x_2, \dots, x_t|\theta')d\mu_a(\theta')}$$

Note that in the right hand side expression, the second term equals to one and the third term is independent from θ. Also $\frac{l(x_2|\theta)}{l(x_1|\theta)}$ is non-decreasing in θ since $\{l(x|\theta)\}_{\theta\in\Theta}$ has MLRP in x and $x_1 < x_2$. So for all $(x_2, \dots, x_t)$ fixed, $\frac{\mu_a(\theta|x_1', x_2, \dots, x_t)}{\mu_a(\theta|x_1, x_2, \dots, x_t)}$ is non-decreasing in θ, and $\{\mu_a(\theta|x_1, x_2, \dots, x_t) : x_1 \in \mathbb{R}\}$ has MLRP in θ. As a result, $\mu_a(\cdot|x_1', x_2, \dots, x_t)$ FOSD $\mu_a(\cdot|x_1, x_2, \dots, x_t)$.

Next we show that for any (μ_a, λ, T)-bandit and first observation $X_1 = x_1$, $V(\mu_a(\cdot|x_1), \lambda, T-1)$ is weakly increasing in x_1. We prove by induction. This holds vacuously when horizon is 1. Suppose it holds for all problems with horizon less than T, we show that it holds for problems with horizon T. Fix $x_1' > x_1$, so

$$\begin{aligned} V(\mu_a(\cdot|x_1'), \lambda, T-1) &= E_{\mu_a(\cdot|x_1')}[X_2 + V(\mu_a(\cdot|X_2, x_1'), \lambda, T-2)] \\ &\geq E_{\mu_a(\cdot|x_1')}[X_2 + V(\mu_a(\cdot|X_2, x_1), \lambda, T-2)] \\ &\geq E_{\mu_a(\cdot|x_1)}[X_2 + V(\mu_a(\cdot|X_2, x_1), \lambda, T-2)] \\ &= V(mu_a(\cdot|x_1), \lambda, T-1) \end{aligned}$$

where the first inequality is by induction hypothesis, and the second inequality is by induction hypothesis and that $P_{\mu_a(\cdot|x_1')}$ FOSD $P_{\mu_a(\cdot|x_1)}$.

Finally, we prove that $\Lambda(C, T) = \Lambda(\mu_{\underline{a}}, T)$. It suffices to show that $V(C, \lambda, T) = V(\mu_{\underline{a}}, \lambda, T)$ for all λ and we prove it by induction on T. The claim is vacuously true when $T = 0$. Suppose it is true for any problem with horizon less than T. We want to show it also holds for horizon T.

$$\begin{aligned} V(C,\lambda,T) &= \inf_{\mu_a \in C} E_{\mu_a}[X_1 + \delta V(C(\cdot|X_1),\lambda,T-1)] \vee \lambda\frac{1-\delta^T}{1-\delta} \\ &= \inf_{\mu_a \in C} E_{\mu_a}[X_1 + \delta V(\mu_{\underline{a}}(\cdot|X_1),\lambda,T-1)] \vee \lambda\frac{1-\delta^T}{1-\delta} \\ &= E_{\mu_{\underline{a}}}[X_1 + \delta V(\mu_{\underline{a}}(\cdot|X_1),\lambda,T-1)] \vee \lambda\frac{1-\delta^T}{1-\delta} \\ &= V(\mu_{\underline{a}},\lambda,T) \end{aligned}$$

where the first and last equality are by recursivity of V, the second equality is by induction hypothesis. To see the third equality, note that $X_1 + \delta V(\mu_{\underline{a}}(\cdot|X_1),\lambda,T-1)$ is a weakly increasing function of X_1 and $P_{\mu_{a_2}}$ FOSD $P_{\mu_{a_1}}$ for all $\underline{a} \le a_1 < a_2 \le \bar{a}$.

The case of $T = \infty$ can be approximated by continuity. □

Proof of Proposition 3.6. The first claim is obvious, since for all $\mu \in C$, $\mu(\cdot|h_t) \in C(\cdot|h_t)$ for all history h_t. By backward construction $U(C)(Z_1,\cdots,Z_T) \le U(\{\mu\})(Z_1,\cdots,Z_T)$ for all $\mu \in C$.

We then prove the second claim. The "only if" statement can be easily proved by contrapositive. We only show the "if" part. For any measurable event $A \subseteq S$, define a payoff process as[1]

$$Z_1 = 1(s_1 \in A),\ Z_2 = \frac{1}{\delta}1(s_2 \in A^c),\ Z_t = 0,\ \forall t = 3,\cdots,T$$

Then for all $\mu \in C$,

$$\begin{aligned} U(\mu)(1(A),\frac{1}{\delta}1(A^c),0,\cdots,0) &= E_\mu[1(s_1 \in A) + \delta\frac{1}{\delta}1(s_2 \in A^c)] \\ &= P_\mu(A) + E_\mu(E_{\mu|s_1}[1(s_2 \in A^c)]) \\ &= P_\mu(A) + P_\mu(A^c) = 1 \end{aligned}$$

thus $\inf_{\mu\in C} U(\mu)(1(A),\frac{1}{\delta}1(A^c),0,\cdots,0) = 1$. We prove the "if" part of the second claim by showing that if $\inf_{\mu\in C} U(\mu)(1(A),\frac{1}{\delta}1(A^c),0,\cdots,0) = U(C)(1(A),\frac{1}{\delta}1(A^c),0,\cdots,0)$ for all measurable $A \subseteq S$, then $\{P_\mu(\cdot) : \mu \in C\}$ has to be singleton.

$$\begin{aligned} &U(C)(1(A),\frac{1}{\delta}1(A^c),0,\cdots,0) \\ &= \inf_{\mu\in C} E_\mu\{1(s_1 \in A) + \delta\frac{1}{\delta}\inf_{\mu'\in C} E_{mu'|s_1}[1(s_2 \in A^c)]\} \\ &= \inf_{\mu\in C} U(\mu)(1(A),\frac{1}{\delta}1(A^c),0,\cdots,0) = 1 \end{aligned}$$

[1]For all measurable event $E \subseteq S$, 1_E is the indicator function that equals to 1 if $s \in E$, and 0 otherwise.

This implies that for arbitrary $\mu \in C$ fixed,

$$P_\mu(A) + E_\mu\{\inf_{\mu' \in C} E_{\mu'|s_1}[1(s_2 \in A^c)]\} \geq 1$$

or equivalently

$$E_\mu\{\inf_{\mu' \in C} E_{\mu'|s_1}[1(s_2 \in A^c)]\} \geq P_\mu(A^c) = E_\mu\{E_{\mu|s_1}[1(s_2 \in A^c)]\}$$

Since $E_{\mu|s_1}[1(s_2 \in A^c)] \geq \inf_{\mu' \in C} E_{\mu'|s_1}[1(s_2 \in A^c)]$ for every s_1, they have to be $P_\mu - a.s.$ equal

$$P_{\mu|s_1}(A^c) = \inf_{\mu' \in C} P_{\mu'|s_1}(A^c),\ P_\mu - a.s.$$

exchange role of A and A^c we have

$$P_{\mu|s_1}(A) = \inf_{\mu' \in C} P_{\mu'|s_1}(A),\ P_\mu - a.s.$$

Since $[P_{\mu|s_1}(A) = 1 - [P_{\mu|s_1}(A^c) = 1 - \inf_{\mu' \in C} P_{\mu'|s_1}(A^c)$,

$$\inf_{\mu' \in C} P_{\mu'|s_1}(A) = \sup_{\mu' \in C} P_{\mu'|s_1}(A),\ P_\mu - a.s.$$

Since this has to hold for all measurable A, we have $P_\mu - a.s.$, $\{P_{\mu'|s_1} : \mu' \in C\}$ is a singleton set. Since μ is arbitrarily chosen from C, we have $\{P_\mu : \mu \in C\}$ is singleton set as well. $\square$