From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2022

Three Essays on Behavioral Economics: Narrow Bracketing and Ambiguity Aversion

Zhang, Mu

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp019c67wq980

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

THREE ESSAYS ON BEHAVIORAL ECONOMICS: NARROW BRACKETING AND AMBIGUITY AVERSION

MU ZHANG

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF ECONOMICS
ADVISER: FARUK GUL

MAY 2022

© Copyright by Mu Zhang, 2022.

All rights reserved.

Abstract

This collection of essays investigates the theory and applications of narrow bracketing and ambiguity aversion.

In Chapter 1, I study two heuristics, narrow bracketing and correlation neglect, that decision makers adopt to simplify the evaluation of risk from multiple sources and axiomatize them as behavioral deviations from the expected utility benchmark by relaxing the independence axiom. When different sources represent different streams of income, my model can explain experimental evidence on violations of first order stochastic dominance and risk aversion over small gambles. When different sources represent consumption in different periods, I show that an Epstein-Zin type utility function can emerge because of narrow bracketing and it can resolve a critique of the original Epstein-Zin model.

In Chapter 2 (co-authored with Rui Tang), we study the implementation problem of a mechanism designer with ambiguity averse agents. The mechanism designer, desiring to implement a choice correspondence, can create ambiguity for agents by committing to multiple allocation rules and transfer schemes without revealing which one to use. By extending the cyclical monotonicity condition from choice functions to choice correspondences, we show that the condition can fully characterize implementable choice correspondences. We then study the implementation problem in supermodular environments and apply our results to public procurement.

In Chapter 3 (co-authored with Pëllumb Reshidi and João Thereze), we study asymptotic learning when the decision-maker is ambiguous about the precision of her information sources. She aims to estimate a state and evaluates outcomes according to the worst-case scenario. Under prior-by-prior updating, ambiguity regarding information sources induces ambiguity about the state. We show this induced ambiguity does not vanish even as the number of information sources grows indefinitely, and characterize the limit set of posteriors. The decision-maker's asymptotic estimate of

the state is generically incorrect. We then consider several applications. For example, we provide a foundation for disagreement among agents with access to the same large dataset.

Acknowledgements

I am deeply indebted to my adviser, Faruk Gul, for his continuous guidance, support, and encouragement throughout my doctoral studies. His invaluable advice helps improve the quality of my research and inspires me to pursue excellence in academics. His generosity has always motivated me to be a better person.

I would like to thank Wolfgang Pesendorfer, Pietro Ortoleva and Xiaosheng Mu for their guidance and insightful suggestions. For many times when I had difficulty making progress on my research, their critical advice helped me find the way out. Their enthusiasm and wisdom for research have shaped my comprehension of economic theory.

I am grateful to Roland Bénabou, Sylvain Chassang, Alessandro Lizzeri, Stephen Morris, Can Urgun, Juan Pablo Xandri Antuna and Leeat Yariv, for their insightful feedback on my research and advice on the job market.

I am thankful to Shaowei Ke, Chen Zhao, Jie Zheng and Jaimie Lien for their help and encouragement. I also would like to thank Rui Tang (co-author of Chapter 2, Pëllumb Reshidi (co-author of Chapter 3) and João Thereze (co-author of Chapter 3) for their enthusiasm and dedication that have been keeping our projects going.

I would like to thank Ellen Graf, Laura Hedden and Erin Casey for providing extraordinary assistance. I thank Princeton University, Dietrich II Economic Theory Center and Stephen Goldfeld Fund for generous financial support.

Many thanks to my wonderful classmates, roommates and friends at Princeton, especially to Zhang Chen, Fei Gao, Lunyang Huang, Dake Li, Lintong Li, Shumiao Ouyang, Peng Wang, Kecy Wu, Haiyue Yu, Qi Zhang, Yi Zhang, and Yinuo Zhang. Their support and company have enriched my years at Princeton.

As always, I want to thank my parents for their unconditional love, support, and encouragement. I am especially indebted to Salome for her company, support and

encouragement, without which I must have been overwhelmed by the stress, anxiety and setbacks during the first five years of my doctoral studies.

To my parents Zhenguang and Qiong.

Contents

Abstract 3

Acknowledgements 5

List of Tables 12

List of Figures 13

1 A Theory of Choice Bracketing under Risk 14

1.1 Introduction 15

1.2 Primitives 23

1.3 Representations 24

1.4 Axioms 28

1.4.1 Narrow Bracketing with Correlation Neglect 28

1.4.2 Narrow Bracketing with Correlation Sensitivity 32

1.4.3 Identification 34

1.5 Proof Sketch 34

1.6 Applications: Portfolio Choices 36

1.6.1 Ranking of Rationality 42

1.6.2 Risk Aversion over Small Gambles 44

1.7 Applications: Time Preferences 45

1.7.1 Epstein-Zin Type Preferences 47

1.7.2 Effects of Heuristics on Optimal Saving 52

1.8 Conclusion 56

2 Maxmin Implementation **58**

2.1 Introduction 59

2.2 Model 67

2.3 Main Result 69

2.4 Multiple Agents 76

2.5 Supermodularity 79

2.6 Application: Allocating a Single Good 84

2.7 Discussion 87

2.7.1 Multiple Agents: Alternative Implementation Concepts 87

2.7.2 Discussion of Assumptions 89

2.7.3 Weak Monotonicity 93

3 Asymptotic Learning with Ambiguous Information **95**

3.1 Introduction 96

3.2 Setup 101

3.3 Asymptotic Beliefs 104

3.3.1 Quadratic Loss with Observable Signals 105

3.3.2 Ambiguity Does Not Vanish 107

3.4 Asymptotic Estimate 110

3.5 Applications 115

3.5.1 Comparative Statics of Ambiguity 115

3.5.2 Misspecification and Ambiguity 116

3.5.3 Aggregating Estimates 117

3.6 Discussion 121

A Appendices for Chapter 1 **124**

A.1 Asymmetric Narrow Bracketing in Infinite Horizon 124

A.2 Proof of Theorem 1.4.1 128

A.3 Proof of Theorem 1.4.2 147
A.4 Proof of Proposition 1.4.3 170
A.5 Proof of Corollary 1.6.2 170
A.6 Proof of Proposition 1.6.3 171
A.7 Proof of Corollary 1.7.1 172
A.8 Proof of Lemma 1.7.5 172
A.9 Proof of Lemma 1.7.6 174
A.10 Proof of Proposition A.1.1 175
A.11 Proof of Corollary A.1.2 176

B Appendices for Chapter 2 **177**

B.1 Revelation Principle 177
B.2 Proof of Theorem 2.3.2 178
B.3 Proof of Theorem 2.4.3 182
B.4 Proof of Proposition 2.5.5 183
B.5 Proof of Proposition 2.5.6 185
B.6 Proof of Theorem 2.6.2 185
B.7 Proof of Proposition 2.7.4 186
B.8 Proof of Proposition 2.7.5 187
B.9 Proof of Proposition 2.7.6 188

C Appendices for Chapter 3 **189**

C.1 Proof of Lemma 3.3.4 189
C.2 Proof of Theorem 3.3.5 190
C.3 Proof of Theorem 3.4.1 199
C.4 Proof of Corollary 3.4.3 203
C.5 Proof of Proposition 3.5.1 204
C.6 Proof of Proposition 3.5.3 205

C.7 Proof of Proposition 3.5.4 206
C.8 Proof of Proposition 3.5.7 207

References **212**

List of Tables

2.1 Expected payments to the firms in the first single-mechanism 62

2.2 Expected payments to the firms in the second single-mechanism . . . 63

List of Figures

1.1 Relationship between models in terms of heuristics. 28

1.2 Rationality ranking of models based on violations of stochastic dominance. 44

1.3 Example of temporal lotteries inducing the same lottery. 48

3.1 Learning From Actions Setup 118

Chapter 1

A Theory of Choice Bracketing under Risk

1.1 Introduction

[1]Decision makers often have to deal with multiple interacting choice problems simultaneously. For example, an investor might need to manage accounts in different financial markets; a worker with both labor income and financial assets has to consider the risk of getting laid off and the risk of a financial crisis, as well as the interaction between them; a consumer might face income risk and interest rate risk in multiple periods; a subject with social preferences might care about the risk in both her own payoff and her opponent's payoff in an experiment. However, aggregating multi-source risk is challenging in practice and decision makers often adopt heuristics to simplify the evaluation process. In this chapter, I focus on two such heuristics, narrow bracketing and correlation neglect.

Narrow bracketing, formalized by Thaler (1985) and Read, Loewenstein, and Rabin (1999), describes a situation where a decision maker (DM) faced with multiple choice problems attempts to solve each in isolation, without taking account of the interactions between the problems. When an agent confronts two choice problems over risky prospects simultaneously, the physical outcome of one choice will affect her ranking of the alternatives in the other problem. A narrow bracketing agent ignores this effect and may end up choosing a stochastically dominated prospect over one that dominates it. While all narrow bracketing agents will choose stochastically dominated prospects in some instances, some forms of narrow bracketing render decision makers more susceptible to violations of stochastic dominance than others. In particular, with the narrow bracketing model that is often used to interpret experimental evidence, it is possible for a decision maker to choose X dollars for sure over $Y > X$ dollars for

[1]This chapter is based on the paper "A Theory of Choice Bracketing under Risk". The paper has been presented in the Twenty-Second ACM Conference on Economics and Computation (EC'21), Decision: Theory, Experiments, and Applications (D-TEA) 2021, Risk, Uncertainty and Decision (RUD) 2021, 2021 North American Summer Meeting of the Econometric Society , 2021 European Summer Meeting of the Econometric Society, 2021 China Meeting of the Econometric Society, Young Economists Symposium (YES) 2021.

sure when these amounts are suitably divided across the two income sources. Building such extreme departures from rationality into agents' behavior might lead to a theory that explains certain anomalies in data at the expense of creating others that are not likely to be present.

Each of the new models of narrow bracketing that I present below allows for a specification that precludes violations of stochastic dominance when there is no uncertainty. Moreover, I provide a ranking of these models based on how likely each one is to generate violations of stochastic dominance; that is, I provide a rationality ranking for narrow bracketing models. I hope that this ranking will be useful for calibrating the level of narrow bracketing and irrationality to specific applications.

The second heuristic, correlation neglect, describes the tendency of agents to ignore the statistical interdependence among different decision problems. When an agent faces two-dimensional risk, the outcome on one dimension will alter the payoff distribution associated with the other dimension. A correlation neglecting agent may underestimate or overestimate the hedging effect that one dimension provides over the other. Such an agent may appear overly risk averse or unwilling to take advantage of diversification opportunities. My analysis reveals that the most extreme form of narrow bracketing, which I call full narrow bracketing, implies correlation neglect. However, in general, it is possible to have a weaker form of narrow bracketing, which I call asymmetric narrow bracketing, with or without correlation neglect. This flexibility is useful because there is plenty of evidence that not all individuals ignore correlation and, in fact, many are averse to it. When evaluating income from two separate sources, this aversion to correlation is an immediate consequence of risk aversion.

I characterize the two above heuristics as behavioral deviations from the expected utility benchmark by relaxing the independence axiom of that model. My axiomatization reveals decision rules that (i) allow for narrow bracketing or correlation neglect

or both and (ii) allow for different degrees of deviations from the benchmark without narrow bracketing. For instance, I allow the DM to violate stochastic dominance in complicated problems but not in simple problems that only involve deterministic options. I show that my asymmetric narrow bracketing model without correlation neglect facilitates a dynamic model that distinguishes between time preference and risk preference without imposing a preference for timing of resolution of uncertainty.

I consider preferences over lotteries with two-dimensional outcome profiles. An outcome profile can be interpreted as the consequences of two decision problems, such as choices over two different classes of risky assets or consumption choices involving two periods. I will call each dimension a *source* of risk and the marginal distribution in some source a *marginal lottery*. My main formal results are two representation theorems. The first builds in correlation neglect and encompasses three cases, (pure) correlation neglect (CN), asymmetric narrow bracketing with correlation neglect (ANB-CN) and full narrow bracketing (FNB). Pure correlation neglect is described by the following type of utility function:

$$V^{CN}(P) = \sum_{x,y} w(x,y) P_1(x) P_2(y)$$

for all $P \in \mathcal{P}$ where P_i is the marginal lottery of P in source i. Hence, the decision maker identifies every P with its marginal distributions and ignores the correlation.

The second case covered by Theorem 1.4.1 is asymmetric narrow bracketing with correlation neglect. Let v be a utility index over one-dimensional gambles and let $CE_v(P_2)$ be the certainty equivalent of the marginal distribution P_2 of P; that is, $CE_v(P_2)$ is the unique x that solves $v(x) = \sum_z v(z) P_2(z)$. Then, the asymmetric narrow bracketing with correlation neglect utility function is:

$$V^{ANB-CN}(P) = \sum_{x} w(x, CE_v(P_2)) P_1(x)$$

for all $P \in \mathcal{P}$. In this case, the decision maker not only ignores the statistical dependence between the two sources of risk but also takes limited account of the physical interaction of the outcomes from two sources. For example, if the two sources are different sources of income, then the utility function above means that the DM is only considering the income effect that arises from the second source "on average".

The last case covered by Theorem 1.4.1 is full narrow bracketing:

$$V^{FNB}(P) = w(CE_u(P_1), CE_v(P_2))$$

for all $P \in \mathcal{P}$. In this case, the DM evaluates both sources of risk in isolation, and therefore, correlation neglect is implied.

Theorem 1.4.1 characterizes the above three types of utility functions with correlation neglect. They follow from two standard assumptions, weak order and monotonicity, weaker versions of standard continuity, independence axioms and the axiom of correlation neutrality.

For my second theorem, I relax the correlation neutrality axiom to a correlation consistency axiom which permits correlation neglect but does not require it. The characterization of Theorem 1.4.2 permits the three representations of Theorem 1.4.1 described above and two other classes of utility functions: the standard, "fully rational", expected utility function

$$V^{EU}(P) = \sum_{x,y} w(x,y)P(x,y)$$

for all $P \in \mathcal{P}$, and the asymmetric narrow bracketing utility function without correlation neglect defined as follows: for any x such that $P(x,y) > 0$ for some y, let $P_{2|x}$ denote the conditional distribution of the outcome in source 2 given x in source 1.

Then, the asymmetric narrow bracketing utility function is

$$V^{ANB}(P) = \sum_x w(x, CE_v(P_{2|x}))P_1(x)$$

for all $P \in \mathcal{P}$. Compared to the asymmetric narrow bracketing with correlation neglect utility function, here the decision maker takes into account the statistical dependence between the two sources of risk. Suppose that the two sources represent consumption levels in two periods. Given each realization of today's consumption, the decision maker evaluates the conditional risk in tomorrow's consumption narrowly using its certainty equivalent. Then she averages over the utility of today's consumption and tomorrow's certainty equivalent. Asymmetric narrow bracketing suggests a heuristic that mimics backward induction in a manner that does not take into account the effect of today's consumption on tomorrow's preferences. This asymmetry reflects the nature of intertemporal problems that future risk is typically more complicated and less accessible than current risk.

Unlike a standard representation theorem that derives a general functional form to include all models, each of my theorems characterizes several seemingly distinct functional forms. In each utility function, the DM either adopts some heuristic, or does not use it at all. This feature follows from my axioms, and is not an ex ante assumption. For instance, intermediate cases such as the weighted average of two utility functions V^{EU} and V^{FNB} are ruled out by my axioms. I argue that such exclusion is consistent with my interpretation of narrow bracketing and correlation neglect as simplifying heuristics, since evaluating a lottery using the weighted average of V^{EU} and V^{FNB} is more complex and demanding than using either V^{EU} or V^{FNB}. Ellis and Freeman (2020) also show that allowing for such intermediate levels of narrow bracketing does not significantly help to explain the data in their experiments.

In Section 1.6 and Section 1.7, I consider applications of my models. First, I interpret the different sources as different streams of income to explain experimental evidence on violations of first order stochastic dominance. I focus on models that satisfy stochastic dominance when there is no uncertainty. I also propose a rationality ranking of these models based on the readiness with which they yield violations of stochastic dominance. This reveals a novel connection between violations of stochastic dominance and deviations from the independence axiom of the expected utility model. Then, I interpret one source as background risk and show that narrow bracketing can explain risk aversion over small gambles.

Finally, I interpret the different sources as consumption in different periods to explore the implications of narrow bracketing and correlation neglect on intertemporal risk attitudes. I show that the full narrow bracketing model is essentially identical to the time preference model studied by Selden (1978) and Selden and Stux (1978). I also show that an Epstein and Zin (1989) type utility function can emerge because of narrow bracketing. This new utility function achieves separation between time and risk preferences without inducing an unreasonably high timing premium as observed by Epstein, Farhi, and Strzalecki (2014). Then, I study how narrow bracketing and correlation neglect affect the optimal saving behavior of a consumer in a simple two-period model. To isolate the effect of narrow bracketing, I consider the standard "certain $\times$ uncertain" setup in the literature of precautionary savings, where there is only one saving decision (in the first period) and income is uncertain in a single period (in the second period). I show that when the consumer's relative risk aversion is larger than the reciprocal of her elasticity of intertemporal substitution, narrow bracketing leads to more risk aversion over future consumption and hence more current saving compared to the expected utility benchmark. To study the effect of correlation neglect, I modify the previous setup by assuming that the consumer knows her future income in the first period and bases her saving decision on this information. Com-

pared to a consumer with correlation neglect, a correlation averse consumer adjusts the current saving levels to compensate for the difference in future income. She saves more when the future income is low and less when the future income is high.

The above applications reveal a novel and possibly surprising connection between the Epstein-Zin type utility function, the most widely used model that separates time and risk preferences in the literature, and violations of first order stochastic dominance, a phenomenon typically considered as an irrational bias. Both the separation of time and risk preferences and violations of stochastic dominance can result from narrow bracketing and be captured by relaxing the independence axiom.

Related Literature. Tversky and Kahneman (1981), Rabin and Weizsäcker (2009) and Ellis and Freeman (2020) provide experimental evidence for narrow bracketing. The experiments in the first two papers isolate narrow bracketing by considering lotteries with independent marginals. Hence, there is no correlation for the DM to ignore. In Ellis and Freeman (2020), each of two assets pays off in only one of two equally likely states, which induces perfect negative correlation between payoffs of the two assets. The authors show that 73% of their subjects are best described by full narrow bracketing, 14% by broad bracketing (expected utility), and only 5% by intermediate cases. However, they do not consider models that separate correlation neglect and narrow bracketing. There is also plenty of experimental evidence for correlation neglect in belief formation (Enke and Zimmermann, 2019), portfolio allocation (Kallir and Sonsino, 2009; Eyster and Weizsäcker, 2016) and school choice (Rees-Jones, Shorrer, and Tergiman, 2020). Those experiments do not address narrow bracketing.

My chapter is related to the growing literature on rationales for narrow bracketing. For instance, Lian (2020) develops an informational theory of narrow bracketing, where the DM cares less about her other decisions when making each decision because of imperfect knowledge of other decisions. Camara (2021) studies the expected util-

ity model with high-dimensional decisions and shows that computational tractability requires the utility index to satisfy a slightly weaker version of additive separability. Complementary to those papers which provide rationales for narrow bracketing, my chapter characterizes the observable implications of narrow bracketing and axiomatizes it as a deviation from expected utility. Moreover, my model can be applied to simple economic settings such as the experiment in Rabin and Weizsäcker (2009), without informational frictions or high-dimensional decisions.

In Vorjohann (2021), each DM is characterized by a broad preference and a narrow preference, both of which can be observed and admit an expected utility representation. Vorjohann (2021) axiomatizes the expected utility index of the narrow preference, which is additively separable across brackets. My work differs from Vorjohann (2021)'s in two aspects. First, the DM in my model has only one preference, which determines whether she is an expected utility maximizer, or adopts some simplifying heuristics. Second, I allow for more flexible forms of narrow bracketing and distinguish it from correlation neglect, while the notion of narrow bracketing in Vorjohann (2021) subsumes correlation neglect.

A strand of recent literature studies the axiomatization of correlation-sensitive models. Ellis and Piccione (2017) model a decision maker who chooses a portfolio of assets among which she may misperceive the joint returns. She has a (possibly misspecified) subjective belief over the correlation of returns and maximizes expected utility. Ellis (2021) further extends the framework and considers a decision maker who deems possible a set of possible correlation structures and evaluates each alternative by the worst expected utility in that set. In my chapter, I only consider the extreme form of correlation misperception, i.e., correlation neglect, since my goal is to separate it from narrow bracketing. I leave the analysis of general correlation misperception for future study. Lanzani (2021) studies a decision maker who is sensitive to the correlation of two risky prospects she needs to compare. By relaxing the transitivity

axiom of the expected utility theory, he provides an axiomatization for the salience theory model in Bordalo, Gennaioli, and Shleifer (2012). My treatment of correlation is complementary to Lanzani (2021). One can interpret that the decision maker in Lanzani (2021) is comparing the *two marginal lotteries* of the same lottery over two-source outcome profiles, while the decision maker in my model is choosing between *two lotteries*.

My application to portfolio choices is related to the recent literature on risk aversion over small gambles. Many remedies have been proposed since the critique of Rabin (2000) on the expected utility theory. However, the results in Barberis, Huang, and Thaler (2006) and Mu, Pomatto, Strack, and Tamuz (2021) suggest that narrow bracketing is necessary to explain commonly observed irregularity in choice under risk. These results motivate my work.

My chapter is also related to the recent work on Epstein-Zin preferences. Epstein, Farhi, and Strzalecki (2014) show that calibrating the Epstein-Zin preferences to macroeconomic and financial data yields a very high early resolution premium. I propose an Epstein-Zin type utility function that emerges from narrow bracketing and addresses this critique.

1.2 Primitives

Let $Z = X_1 \times X_2$, where X_i is the set of outcomes in *source* $i \in \{1, 2\}$. I assume that X_i is a nontrivial compact (closed and bounded) interval on the real line. I call $(x_1, x_2) \in Z$ an *outcome profile* and x_i is the outcome in source i for $i \in \{1, 2\}$.

A *(joint) lottery*, P, is a probability distribution over Z with a finite support. I endow the set of all lotteries, $\mathcal{P} = \mathcal{L}^0(Z)$, with the topology of weak convergence and the standard mixture operation. For $P \in \mathcal{P}$, the *marginal lottery* of P in source 1 is denoted $P_1 \in \mathcal{L}^0(X_1)$. That is, $P_1(x) = \sum_{y \in X_2} P(x, y)$ for all $x \in X_1$. The marginal

lottery P_2 in source 2 is defined similarly. Let $\hat{\mathcal{P}} = \mathcal{L}^0(X_1) \times \mathcal{L}^0(X_2) \subset \mathcal{P}$ denote the set of *product lotteries*. For $P \in \mathcal{P}$, the pair $(P_1, P_2) \in \mathcal{P}$ is a product lottery with the same marginals as P. When there is no risk of confusion, I use (x_1, x_2) to denote the degenerate lottery that yields the outcome profile $(x_1, x_2) \in Z$ for sure and interpret Z as a subset of $\mathcal{P}$.

The primitive of my analysis is a binary relation $\succsim$ on $\mathcal{P}$.

1.3 Representations

In this section, I introduce different decision rules adopted by a decision maker faced with two-source risk. I start with the expected utility model as the benchmark. A binary relation $\succsim$ is said to admit an *expected utility* (EU) representation if it can be represented by V^{EU} with

$$V^{EU}(P) = \sum_{x,y} w(x,y)P(x,y)$$

for all $P \in \mathcal{P}$, where w is a continuous and strictly monotonic utility index on Z. In Section 1.6 and Section 1.7, I will study specifications of w when outcomes represent income streams from two sources or consumption in two periods.

As discussed in the introduction, people in practice usually deviate from the benchmark systematically by adopting some simplifying heuristics. In this chapter, I will focus on two such heuristics: choice bracketing and correlation neglect. I say that a binary relation $\succsim$ admits a *full narrow bracketing* (FNB) representation if it is represented by V^{FNB}:

$$V^{FNB}(P) = w(CE_u(P_1), CE_v(P_2))$$

for all $P \in \mathcal{P}$, where w is a continuous and strictly monotonic utility index on Z and u, v are continuous and strictly monotonic utility indices on X_1, X_2 respectively. $CE_u(P_1)$ is the certainty equivalent of the marginal distribution P_1 of P in source 1 under utility index u; that is, $CE_u(P_1)$ is the unique x that solves $u(x) = \sum_z u(z)P_1(z)$. $CE_v(P_2)$ is defined similarly. Hence, a full narrow bracketing decision maker evaluates each marginal lottery, P_i, independently and aggregates the two evaluations to determine the overall utility of lottery P. It follows that she also ignores the correlation between the marginals.

The full narrow bracketing model might be too restrictive as it necessitates correlation neglect. Recent experimental literature has shown the prevalence of correlation non-neutrality in different economic scenarios. For instance, Andersen, Harrison, Lau, and Rutström (2018) find strong evidence for correlation aversion over lotteries with intertemporal income profiles, that is, subjects are averse to positive correlation between income in different periods. Lanier, Miao, Quah, and Zhong (2020) confirm this finding using non-parametric analysis and a different experimental design. For atemporal settings, Ebert and van de Kuilen (2016) and Attema, L'haridon, and van de Kuilen (2019) find correlation non-neutrality for two-source risk in monetary outcomes and outcomes in terms of waiting time or longevity. Thus, it is worthwhile to develop a model that separates choice bracketing from correlation neglect.

Now I introduce the asymmetry narrow bracketing model, where the decision maker only narrowly bracket risk in one source, instead of both sources. Without loss of generality, I denote by source 2 the source of narrow bracketing. This labeling makes sense in practice when the risk in source 2 is more complicated and/or less accessible than the risk in source 1. For instance, I interpret the outcome in source 1 as today's consumption and the outcome in source 2 as tomorrow's consumption, because evaluating tomorrow's risk requires the decision maker to assess the hypothetical circumstances she is not facing at present. Future risk is also more complex since the

distribution over consumption paths is usually presented in the form of a probability tree starting with the current consumption. Similarly, in social preferences, I interpret the outcomes in source 1 and source 2 as the payoffs of the decision maker and her opponent respectively, because it is more complicated for the decision maker to put herself in the opponent's shoes. In the case with background risk, it is less clear whether or not the decision maker is more inclined to narrowly bracket the background risk. One might argue that outcomes generated by the prospect the decision maker is facing now are more accessible than background wealth and hence source 2 should represent the background risk. However, since the background risk is fixed and the decision maker is choosing over different risky prospects at hand, she might evaluate a risky prospect conditional on the realization of background wealth. In this case, source 1 should represent the background risk.

For each $P \in \mathcal{P}$ and $x \in X_1$ such that $P(x, y) > 0$ for some $y \in X_2$, let $P_{2|x}$ denote the conditional distribution of the outcome in source 2 given x in source 1. A binary relation $\succsim$ admits an *asymmetric narrow bracketing* (ANB) representation if it is represented by V^{ANB}:

$$V^{ANB}(P) = \sum_x w(x, CE_v(P_{2|x}))P_1(x)$$

for all $P \in \mathcal{P}$, where w, v are continuous and strictly monotonic utility indices on Z, X_2 respectively. To assesses the overall utility of lottery P, a decision maker with asymmetric narrow bracketing first replaces each conditional lottery in source 2 with its certainty equivalent and then evaluates the new lottery using expected utility. The decision maker narrowly brackets the risk in source 2 but not in source 1. Moreover, she takes into account the correlation between outcomes in two sources. Since the notion of correlation aversion typically involves lotteries with degenerate conditional lotteries in source 2 (e.g., Bommier (2007)), V^{ANB} and V^{EU} share the same attitude

towards correlation if they have the same utility index w over deterministic outcome profiles.

When each outcome profile represeents a consumption path, the asymmetric narrow bracketing model combines narrow bracketing with backward induction and resembles the functional form of recursive preferences such as Epstein and Zin (1989). In Section 1.7, I show that an Epstein-Zin type utility function can arise due to narrow bracketing instead of a preference for timing of resolution of uncertainty.

Now I consider a decision maker with correlation neglect and no narrow bracketing. A binary relation $\succsim$ admits an *(pure) correlation neglect* (CN) representation if it is represented by V^{CN}:

$$V^{CN}(P) = \sum_{x,y} w(x,y)P_1(x)P_2(y)$$

for all $P \in \mathcal{P}$, where w is a continuous and strictly monotonic utility index on Z. A decision maker with correlation neglect rationally aggregates independent risk from two sources, but fails to appreciate the correlation.

My final utility function combines asymmetric narrow bracketing with correlation neglect (ANB-CN), and is described below:

$$V^{ANB-CN}(P) = \sum_{x} w(x, CE_v(P_2))P_1(x)$$

for all $P \in \mathcal{P}$, where w, v are continuous and strictly monotonic utility indices on Z, X_2 respectively. Compared to V^{ANB}, here the decision maker ignores the statistical dependence between the two sources of risk: she first determines the certainty equivalent of the marginal lottery of P in source 2 and then assesses the overall utility of P as if it yields that certainty equivalent in source 2 for sure.

Figure 1.1 depicts the heuristics involved in the five models introduced above. Given the expected utility benchmark, V^{CN} exhibits correlation neglect and no narrow bracketing while V^{ANB} exhibits narrow bracketing in source 2 and no correlation

neglect. V^{ANB-CN} combines the previous two heuristics, and V^{FNB} further involves narrow bracketing in source 1.

Figure 1.1: Relationship between models in terms of heuristics.

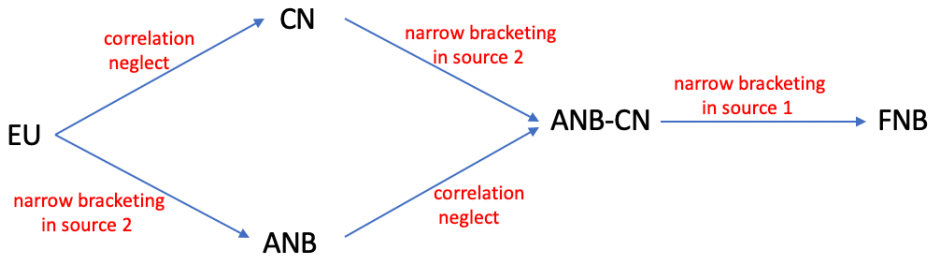

1.4 Axioms

I will present two representation theorems. My Theorem 1.4.1 identifies axioms that ensure $\succsim$ can be represented by one of these three utility functions with correlation neglect: V^{CN}, V^{ANB-CN} or V^{FNB}. Moreover, any preference that can be represented by one of these utility functions must satisfy the axioms in Theorem 1.4.1. In Theorem 1.4.2, I weaken correlation neglect to incorporate the other two utility functions: the expected utility benchmark V^{EU} and asymmetric narrow bracketing V^{ANB}.

1.4.1 Narrow Bracketing with Correlation Neglect

In this section, I present my axioms for representations that exhibit correlation neglect. The first of four common axioms shared by my two main theorems is rationality.

Weak Order: $\succsim$ is complete and transitive.

The next axiom is a version of monotonicity for my two-dimensional setting. With single-source risk, there is an agreed notion first order stochastic dominance. However,

it is possible to extend this notion to a setting with multiple sources in several ways. I will choose a weak notion of dominance, which only involves comparing each lottery P with degenerate lotteries $(x, y) \in Z$. For any marginal lottery $p \in \mathcal{L}^0(X_1) \cup \mathcal{L}^0(X_2)$, let $\text{supp}(p) := \{x \in X_1 \cup X_2 : p(x) > 0\}$ be the support of p. For any $P, (x, y) \in \mathcal{P}$, I say P *dominates* (x, y) if $P \neq (x, y)$ and $x' \geq x, y' \geq y$ for all $x' \in \text{supp}(P_1), y' \in \text{supp}(P_2)$. Symmetrically, I say (x, y) *dominates* P if $P \neq (x, y)$ and $x' \leq x, y' \leq y$ for all $x' \in \text{supp}(P_1), y' \in \text{supp}(P_2)$. Axiom 2 states that the preference $\succsim$ is monotonic with respect to this notion of dominance.

Monotonicity: For $P \in \mathcal{P}$ and $(x, y) \in Z$, $P \succ (x, y)$ if P dominates (x, y) and $(x, y) \succ P$ if (x, y) dominates P.

My continuity axiom combines three weaker notions of the standard continuity axiom: Continuity 1 is mixture space continuity (i.e., continuity in probabilities); Continuity 2 is continuity in outcomes while Continuity 3 combines the two but is restricted to product lotteries.

Continuity 1: For $P, R, Q \in \mathcal{P}$, the sets $\{\alpha \in [0, 1] : \alpha P + (1 - \alpha)Q \succ R\}$ and $\{\alpha \in [0, 1] : R \succ \alpha P + (1 - \alpha)Q\}$ are open subsets of $[0, 1]$.

Continuity 2, below, asserts that if increasing (decreasing) every outcome in the support of P by $(\varepsilon^1, \varepsilon^2)$ renders P better (worse) than Q, for any positive $\varepsilon^1, \varepsilon^2$, then P must also be better (worse) than Q. Since the set Z is bounded, I need to deal with the possibility that increasing some of the prizes in the support of P may not be feasible. Formally, for any $P \in \mathcal{P}$, choose $\eta > 0$ small enough so that $P(x, y) \cdot P(x', y') > 0$, $|x - x'| \leq \eta$ and $|y - y'| \leq \eta$ implies $(x, y) = (x', y')$. Then, for $\varepsilon = (\varepsilon^1, \varepsilon^2)$ such that $\varepsilon^1, \varepsilon^2 \in (-\eta, \eta)$, define $\phi : Z \to Z$ as follows: $\phi_\varepsilon(x, y) = (x + \varepsilon^1, y + \varepsilon^2)$ if $(x + \varepsilon^1, y + \varepsilon^2) \in Z$; otherwise, $\phi_\varepsilon(x, y) = (x', y')$ such that (x', y') is the element of Z closest to (x, y) with respect to the distance $d((x, y), (x', y')) = |x - x'| + |y - y'|$. Since I have chosen $\varepsilon^1, \varepsilon^2$ sufficiently small, the restriction of ϕ to the support of P

is one-to-one. Then, define $P_\varepsilon \in \mathcal{P}$ as follows

$$P_\varepsilon(\phi_\varepsilon(x,y)) = P(x,y)$$

if $P(x,y) > 0$ and set $P_\varepsilon(x',y') = 0$ if $(x',y') \neq \phi_\varepsilon(x,y)$ for any (x,y) with $P(x,y) > 0$.

Continuity 2: For all $P, Q \in \mathcal{P}$ and ε_n converging to $(0,0)$,

$$P_{\varepsilon_n} \succsim Q \text{ for all } n \text{ implies } P \succsim Q \text{ and } Q \succsim P_{\varepsilon_n}, \text{ for all } n \text{ implies } Q \succsim P.$$

For product lotteries, I will assume strong continuity:

Continuity 3: For all $Q \in \mathcal{P}$, the sets $\{P \in \hat{\mathcal{P}} : P \succ Q\}$ and $\{P \in \hat{\mathcal{P}} : Q \succ P\}$ are open subsets of $\hat{\mathcal{P}}$.

I will refer to the conjunction of the three notions above as **Continuity**. Note that Continuity is weaker than the strong continuity axiom of expected utility theory with monetary prizes, which is Continuity 3 by replacing $\hat{\mathcal{P}}$ with $\mathcal{P}$. Note that Continuity 1 and Continuity 2 would be implied by this strong continuity. However, Continuity does not imply strong continuity.

Below, I consider an alternative to the von Neumann-Morgenstern (vNM) independence axiom. Recall that expected utility theory preferences satisfy the following stronger version of the independence axiom: $P \succ Q$, $R \sim S$ implies $\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$ for $\alpha \in (0,1)$. In my alternative imposed on product lotteries $\hat{\mathcal{P}}$, I restrict the condition above to three situations in which P-R, Q-S or R-S are sufficiently similar. In order to guarantee that the mixture of two product lotteries remain a product lottery, I will require that P and R are identical in some source and so are Q and S.

I define the *conditional preference in source 1* given marginal lottery $q \in \mathcal{L}^0(X_2)$ in source 2, $\succsim_{1|q}$, as the restriction of $\succsim$ on $\mathcal{L}^0(X_1) \times \{q\}$, that is, $p \succsim_{1|q} p'$ if and only if $(p, q) \succsim (p', q)$ for each $p, p' \in \mathcal{L}^0(X_1)$. Symmetrically, for each $p \in \mathcal{L}^0(X_1)$, $\succsim_{2|p}$ denotes the conditional preference in source 2 given p in source 1. For $i \in \{1, 2\}$, I denote $-i \neq i$ and $-i \in \{1, 2\}$. Then I say P and Q are *conditionally indifferent in source* i if there exists $z \in X_{-i}$ such that $P_i \sim_{i|z} Q_i$. Conditional indifference captures the notion of sufficient similarity of marginal lotteries in one source.

The first notion of sufficient similarity requires R to be the same as S, and P, Q and R to have the same marginal lottery in one source. The second notion requires P to be conditionally indifferent to R in one source and identical to R in the other; similarly, Q and S need to be conditionally indifferent in one source and identical in the other. The third notion requires that R is the same as S, identical to P in one source and identical to Q in the other, and that P and Q are conditionally indifferent in source 2.

Independence*: For $P, Q, R, S \in \hat{\mathcal{P}}$ and $\alpha \in (0, 1)$,

$$P \succ Q \Longrightarrow \alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S \quad \text{if}$$

(i) $R = S$ and for some $i \in \{1, 2\}$, $P_i = Q_i = R_i$ or

(ii) $R \sim S$ and for some $i, j \in \{1, 2\}$, $x \in X_i$ and $y \in X_j$, $P_i = R_i$, $Q_j = S_j$, $P_{-i} \sim_{-i|x} R_{-i}$ and $Q_{-j} \sim_{-j|y} S_{-j}$ or

(iii) $R = S = (P_1, Q_2)$ or (Q_1, P_2) and for some $x \in X_1$, $P_2 \sim_{2|x} Q_2$.

Now I briefly discuss the three parts in Independence*. Part (i) guarantees that given a fixed marginal lottery in one source, the conditional preference in the other must admit an expected utility representation. Hence, I maintain the independence axiom within each source and only relax inter-source independence. Part (ii) becomes

a trivial implication of part (i) if P is indifferent to R or Q is indifferent to S. Suppose that these two conditions do not hold, then the requirement that P and R are conditionally indifferent in source $-i$ implies that how the DM ranks P_{-i} and R_{-i} depends on the marginal risk in the other source since $P_{-i} \not\sim_{-i|P_i} R_{-i}$ and $P_{-i} \sim_{-i|x} R_{-i}$ for some $x \in X_2$. This implies that the DM does not *narrowly bracket* P_{-i} and R_{-i} in source $-i$. Similarly, the DM does not narrowly bracket Q_{-j} and S_{-j} in source $-j$. Thus, part (ii) essentially states that the independence axiom holds when the DM does not narrowly bracket marginal lotteries in P-R and Q-S. Part (iii) strengthens part (ii) in the case where R is the same as S, identical to P in one source and identical to Q in the other. In this case, the independence axiom holds if the DM does not narrowly bracket the P_2 and Q_2. Part (iii) is important to rule out a symmetric counterpart of V^{ANB-CN} in Theorem 1.4.1 where the roles of two sources are reversed.

My final axiom asserts that the DM identifies every lottery with the product of its marginals; that is, she ignores correlation.

Correlation Neutrality: For all $P \in \mathcal{P}$, $P \sim (P_1, P_2)$.

Now I am ready to state my first representation theorem under correlation neglect.

Theorem 1.4.1 *A binary relation on $\mathcal{P}$ satisfies Weak Order, Monotonicity, Continuity, Independence* and Correlation Neutrality if and only if it can be represented by one of the following types of utility functions: V^{CN}, V^{ANB-CN} or V^{FNB}.*

1.4.2 Narrow Bracketing with Correlation Sensitivity

The following axiom relaxes Correlation Neutrality to incorporate representations that are sensitive to the correlation structure of risk in different sources, that is, the expected utility benchmark V^{EU} and asymmetric narrow bracketing V^{ANB}.

Correlation Consistency: For $P, Q, R, S \in \mathcal{P}$ and $\alpha \in (0,1)$,

$$\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S \quad \text{if}$$

$P \succ Q$, $R \sim S$, for $i = 1, 2$, $P_i = Q_i$, $supp(P_1) \cap supp(R_1) = supp(Q_1) \cap supp(S_1) = \emptyset$.

To understand Correlation Consistency, note that since P and Q have the same marginal lotteries, $P \succ Q$ means that the DM prefers the correlation structure of P to that of Q. Then, the axiom states that this preference is preserved when R is indifferent to S, and the mixture of P and R, Q and S does not affect the correlation structures of P and Q. Here, I interpret the correlation structure of P as the profile of conditional lotteries $\{P_{2|x}\}_{x \in \text{supp}(P_1)}$ in source 2, i.e., the source that the DM is more inclined to narrowly bracket.

My Theorem 1.4.2 states that, by relaxing Correlation Neutrality in Theorem 1.4.1 to Correlation Consistency, I allow for two more utility functions V^{EU} and V^{ANB}.

Theorem 1.4.2 *A binary relation on $\mathcal{P}$ satisfies Weak Order, Monotonicity, Continuity, Independence* and Correlation Consistency if and only if it can be represented by one of the following types of utility functions: V^{EU}, V^{ANB}, V^{CN}, V^{ANB-CN} or V^{FNB}.*

Theorem 1.4.1 and Theorem 1.4.2 characterize behavioral decision rules differing in the involved behavioral heuristics: correlation neglect, narrow bracketing in source 1 and narrow bracketing in source 2. In each representation, a heuristic is either adopted or completely ruled out. This feature results from my axioms, instead of an ex ante assumption, and is in contrast to a standard representation theorem in the literature, which axiomatizes a general functional form that incorporates all models.

1.4.3 Identification

The identification of utility functions characterized in Theorem 1.4.1 and Theorem 1.4.2 is clear: except for the utility index over deterministic outcome profiles w in V^{FNB}, all other utility indices are unique up to a positive affine transformation.

Proposition 1.4.3 *i). If $\succsim$ is represented by V^{EU} or V^{CN} with indices w and w', then there exist $a > 0$ and $b \in \mathbb{R}$ such that $w'(x,y) = aw(x,y) + b$ for all $(x,y) \in Z$.*
ii). If $\succsim$ is represented by V^{ANB} or V^{ANB-CN} with indices (w,v) and (w',v'), then there exist $a, a' > 0$ and $b, b' \in \mathbb{R}$ such that $w'(x,y) = aw(x,y) + b$ for all $(x,y) \in Z$ and $v'(y) = a'v(y) + b'$ for all $y \in X_2$.
iii). If $\succsim$ is represented by V^{FNB} with indices (w,u,v) and (w',u',v'), then there exist $a, a' > 0$ and $b, b' \in \mathbb{R}$ and $\phi : w(Z) \to \mathbb{R}$ continuous and strictly increasing such that $w'(x,y) = \phi(w(x,y))$ for all $(x,y) \in Z$, $u'(x) = au(x) + b$ for all $x \in X_1$ and $v'(y) = a'v(y) + b'$ for all $y \in X_2$.

1.5 Proof Sketch

In this section, I briefly sketch the proofs of Theorem 1.4.1 and Theorem 1.4.2. I start with Theorem 1.4.1, the representation theorem for models that exhibit correlation neglect. It will serve as an important intermediate result for the proof of Theorem 1.4.2.

By Correlation Neutrality, it suffices to focus on the restriction of $\succsim$ to the set of product lotteries $\hat{\mathcal{P}}$. Recall that $\succsim_{i|q}$ is the conditional preference over marginal lotteries in source i given q in the other source. By Continuity and part (i) of Independence*, $\succsim_{i|q}$ admits an expected utility representation.

If $\succsim_{i|q}$ is independent of $q \in \mathcal{L}^0(X_{-i})$ for each $i = 1, 2$, then the decision maker adopts narrow bracketing in both sources and evaluates the marginal distributions in isolation. Suppose $\succsim_i := \succsim_{i|q}$ is represented by the expected utility index v_i. Continu-

ity and Monotonicity guarantee that v_i is continuous and strictly increasing and the certainty equivalent function CE_{v_i} is well defined. Hence, the decision maker is indifferent between $P \in \mathcal{P}$ and $(CE_{v_1}(P_1), CE_{v_2}(P_2))$, and her preference is represented by some V^{FNB}.

If $\succsim_{2|q}$ is independent of $q \in \mathcal{L}^0(X_1)$, and $\succsim_{1|p} \neq \succsim_{1|p'}$ for some $p, p' \in \mathcal{L}^0(X_2)$, then the decision maker adopts narrow bracketing in source 2 but not in source 1. I can focus on the restriction of $\succsim$ to $\mathcal{L}^0(X_1) \times X_2$ and the goal is to show:

$$P \succsim Q \Longleftrightarrow \alpha P + (1-\alpha)R \succsim \alpha Q + (1-\alpha)S, \ \forall\ \alpha \in (0,1), \tag{1.1}$$

if $R \sim S$ and $P_2 = R_2 = y_1, Q_2 = S_2 = y_2$ with $y_1, y_2 \in X_2$. This property guarantees that the decision maker admits a utility function which is linear in the marginal distributions of the outcome in source 1. Then one can show that it is exactly the V^{ANB-CN} representation.

If $\succsim_{2|q} \neq \succsim_{2|q'}$ for some $q, q' \in \mathcal{L}^0(X_1)$, then the decision maker does not adopt narrow bracketing in source 2. Part (ii) of Independence* implies a local version of Property 1.1; that is, it holds under an additional condition that P and R, Q and S are conditionally indifferent in source 1. This requires the existence of $y_1', y_2' \in X_2$ with $P_1 \sim_{1|y_1'} R_1$ and $Q_1 \sim_{1|y_2'} S_1$ when $P \not\sim R$ and $Q \not\sim S$. Then the technical part of the proof is to bridge the gap between the local version and the global version of Property 1.1. Intuitively, if the decision maker never exhibits narrow bracketing in source 1 locally, then she does not exhibit narrow bracketing in source 1 globally. By part (iii) of Independence*, which states that the decision maker is more inclined to narrowly bracket risk in source 2, I show that the decicion maker's preference must be represented by some V^{CN}. This completes the proof for Theorem 1.4.1.

The proof of Theorem 1.4.2 can be decomposed into four steps. First, I define $\hat{\succsim}$ on $\mathcal{P}$ which satisfies Correlation Neutrality and agrees with $\succsim$ on $\hat{\mathcal{P}}$. Easy to very

that $\hat{\succsim}$ satisfies the axioms in Theorem 1.4.1 and hence can be represented by either V^{CN} or V^{ANB-CN} or V^{FNB}. If Correlation Neutrality holds, then $\hat{\succsim} = \succsim$. From now on, suppose that Correlation Neutrality fails. I say two lotteries P and Q are compatible if their marginal distributions in source 1 have disjoint supports, that is, $P_1(x) \cdot Q_1(x) = 0$ for all $x \in X_1$. Compatibility guarantees that the mixture between P and Q does not alter their correlation structure in terms of conditional distributions of outcomes in source 2. The second step is to show that the standard independence axiom holds subject to this compatibility constraint. Third, I extend $\succsim$ to a richer mixture space that mimics the space of temporal lotteries in Kreps and Porteus (1978). I show that the extended preference satisfies the standard mixture space axioms and admits a linear representation. Finally, exploiting the consistency of the two representations of $\succsim$ restricted to $\hat{\mathcal{P}}$ in the previous steps, I conclude that $\succsim$ can be represented by either V^{EU} or V^{ANB} when Correlation Neutrality fails.

1.6 Applications: Portfolio Choices

In this section, I consider a DM faced with uncertainty in income streams from multiple sources. I show that narrow bracketing can explain violations of first order stochastic dominance observed in experiments, and risk aversion over small gambles with background risk. I also develop a rationality ranking of utility functions axiomatized in Theorem 1.4.2 based on the readiness with which they yield violations of stochastic dominance. Consistent with the intuition, the more behavior heuristics a utility function encompasses, the more likely it is to yield violations of stochastic dominance and hence the more "irrational" it is.

The analysis in this section can accommodate various economic settings. For instance, a DM who contemplates the optimal portfolio choice of financial assets needs to take into account the expected returns and volatility of different assets as well as

the comovement of their prices; a DM with both labor income and capital income has to consider the probability of getting laid off, the probability of a financial crisis, and the probability of both happening simultaneously; a DM who decides whether to purchase small-scale insurance for her brand new smart phone should bear in mind the background risk of lifetime income and how it might correlate with the risk associated with the smart phone. Unless otherwise specified, I will keep the terminology of portfolio choices.

I assume that the outcome space in each income source is a compact interval $X_1 = X_2 = X := [-\bar{x}, \bar{x}] \in \mathbb{R}$ with $\bar{x} > 0$ sufficiently large. The outcome $x_i \in X$ in each source can be interpreted as the monetary income in source $i \in \{1, 2\}$, and negative values of x_i represent losses.

Recall that $\succsim_{1|r}$ denotes the conditional preference in source 1 given marginal lottery $r \in \mathcal{L}^0(X)$ in source 2. When $r = \delta_0$, I call it the *narrow preference in source 1* and denote it by $\succsim_1$, that is, $p \succsim_1 q$ if and only if $(p, 0) \succsim (q, 0)$. Since the marginal lottery in source 2 is fixed at δ_0, the comparison of lotteries $(p, 0)$ and $(q, 0)$ can be interpreted as the comparison of p and q in source 1 as if source 2 does not exist. Symmetrically, $\succsim_2$ denotes the narrow preference in source 2. I assume that narrow preferences in two sources are identical since both of them represent the decision maker's preference over risky prospects of money.

Symmetry: $\succsim_1 = \succsim_2$.

Along with Monotonicity, Continuity and part (i) of Independence*, Symmetry implies that the two narrow preferences $\succsim_1$ and $\succsim_2$ admit the same expected utility representation with index $u : X \to \mathbb{R}$ continuous and strictly increasing.

Now I formally introduce the notion of first order stochastic dominance. If the DM's income streams in two sources are determined by the lottery $P \in \mathcal{P}$, then the distribution over her final wealth can be denoted by $f[P] \in \mathcal{L}^0(\mathbb{R})$ where for each

$z \in \mathbb{R}$,

$$f[P](z) = \sum_{(x,y):x+y=z} P(x,y).$$

For any two distributions $p, q \in \mathcal{L}^0(\mathbb{R})$, I say that p first order stochastically dominates q, denoted by $p \succ_{FOSD} q$, if $p \neq q$ and $\sum_{x \leq z} q(x) \geq \sum_{x \leq z} p(x)$ for all $z \in \mathbb{R}$. Then I define that $\succsim$ satisfies *(first order) stochastic dominance* if and only if for each $P, Q \in \mathcal{P}$, $f[P] \succ_{FOSD} f[Q]$ implies $P \succ Q$ and $f[P] = f[Q]$ implies $P \sim Q$.

Intuitively, a rational DM cares about final wealth and prefers more money than less, and hence would choose P over Q if $f[P] \succ_{FOSD} f[Q]$ since P assigns higher probability to more final wealth. This suggests that the preference of a rational DM should satisfy stochastic dominance. However, experimental evidence shows that many subjects violate stochastic dominance. Consider the following experiment studied in Tversky and Kahneman (1981) and Rabin and Weizsäcker (2009).

Example 1.6.1 Suppose you face the following pair of concurrent decisions. All lotteries are independent. First examine both decisions, then indicate ymy choices. Both choices will be payoff relevant, i.e., the gains and losses will be added to ymy overall payment.

Decision 1: Choose between:

A. A sure gain of $2.40.

B. A 25 percent chance to gain $10.00, and a 75 percent chance to gain $0.

Decision 2: Choose between:

C. A sure loss of $7.50.

D. A 75 percent chance to lose $10.00, and a 25 percent chance to lose $0.

Across different treatments in Tversky and Kahneman (1981) and Rabin and Weizsäcker (2009), a large fraction of subjects choose A in decision 1 and D in decision

2. Those choices violate stochastic dominance because

$$f[(B,C)] = \frac{3}{4}\delta_{-7.50} + \frac{3}{4}\delta_{2.50} \succ_{FOSD} \frac{3}{4}\delta_{-7.60} + \frac{3}{4}\delta_{2.40} = f[(A,D)].$$

Such violations are stark since the lottery over final wealth induced by B and C is that induced by A and D plus a sure gain of \$0.10.

Like other experiments on violations of stochastic dominance, Example 1.6.1 involves risk in at least some of the options. This design makes sense because when all options are deterministic, the decision problem is so simple that one can confidently believe that almost all subjects should be able to sum up the monetary prizes across different sources and choose the option delivering the highest final wealth. Hence, I assume that the DM does not violate dominance when there is no risk.

Stochastic Dominance without Risk: For each $x, y, x', y' \in X$, $(x,y) \succsim (x',y')$ if and only if $x + y \geq x' + y'$.

Then I have the following direct corollary of Theorem 1.4.2.

Corollary 1.6.2 *Suppose that $\succsim$ satisfies the axioms stated in Theorem 1.4.2 and can be represented by one of V^{EU}, V^{ANB}, V^{CN}, V^{ANB-CN} or V^{FNB}. Then $\succsim$ satisfies Symmetry and Stochastic Dominance without Risk if and only if those utility functions can be rewritten such that for each $P \in \mathcal{P}$,*

$$\begin{aligned}
V^{EU}(P) &= \sum_{x,y} u(x+y)P(x,y);\\
V^{CN}(P) &= \sum_{x,y} u(x+y)P_1(x)P_2(y);\\
V^{ANB}(P) &= \sum_{x} u(x + CE_u(P_{2|x}))P_1(x);\\
V^{ANB-CN}(P) &= \sum_{x} u(x + CE_u(P_2))P_1(x);\\
V^{FNB}(P) &= u(CE_u(P_1) + CE_u(P_2)).
\end{aligned}$$

where $u : [-2\bar{x}, 2\bar{x}] \to \mathbb{R}$ *is continuous and strictly monotonic.*

Since u is strictly monotonic, a preference that is represented by V^{FNB} can also be represented by $CE_u(P_1) + CE_u(P_2)$. Hence, according to Corollary 1.6.2, a DM with full narrow bracketing evaluates a lottery over two income sources by summing up the *certainty equivalents* of the marginal lotteries in each source. In contrast, an alternative utility function of full narrow bracketing, which has been widely applied in the behavioral and experimental literature, entails the summation of *expected utilities* of the marginal lotteries. Concretely, the utility of $P \in \mathcal{P}$ is

$$\hat{V}^{FNB}(P) = \sum_x u(x)P_1(x) + \sum_y u(y)P_2(y).$$

$\hat{V}^{FNB}$ can violate stochastic dominance even without risk. To see this intuitively, suppose the DM is choosing between a portfolio P which delivers \$1 in both assets, with a portfolio Q which delivers \$2 in asset one and \$0 in asset two. If the DM is risk averse, then she will strictly prefer portfolio P to portfolio Q, although holding either portfolio delivers total income \$2 for sure. Hence, V^{FNB} should be more reasonable than $\hat{V}^{FNB}$ in portfolio choice problems. For most applications in the literature, replacing $\hat{V}^{FNB}$ with V^{FNB} will not significantly alter the qualitative results.

I argue that V^{FNB} can explain violations of stochastic dominance in Example 1.6.1. To see this, suppose that the subject's preference is represented by V^{FNB} with

$$u(x) = \begin{cases} \sqrt{x}, & \text{if } x \geq 0, x \in X, \\ -2\sqrt{-x}, & \text{if } x < 0, x \in X. \end{cases} \tag{1.2}$$

u is a gain-loss utility index with CRRA risk preference and loss aversion parameter 2, which is a standard parametrization in the behavioral economics literature. Then

the subject will simultaneously choose A in decision 1 and D in decision 2 since

$$CE_v(A) = 2.4 > 0.625 = CE_v(B), CE_v(D) = -5.625 > -7.5 = CE_v(C).$$

Since lotteries are assumed to be independent from each other, with a similar parametrization of u, violations of stochastic dominance can also be explained by V^{ANB} and V^{ANB-CN}.

Before moving on, I briefly discuss the *partial narrow bracketing* model studied in Barberis, Huang, and Thaler (2006) and Rabin and Weizsäcker (2009) and Ellis and Freeman (2020). Without violating Stochastic Dominance without Risk, the model can be adapted to my setting as follows:

$$\begin{aligned} V^{PNB}(P) &= \alpha V^{EU}(P) + (1-\alpha)V^{FNB} \\ &= \alpha \sum_{x,y} u(x+y)P(x,y) + (1-\alpha)u(CE_u(P_1) + CE_u(P_2)), \end{aligned}$$

for each $P \in \mathcal{P}$ where $\alpha \in [0,1]$ and $1-\alpha$ determines the degree of narrow bracketing.

The partial narrow bracketing utility function V^{PNB} does not satisfy the axioms in Theorem 1.4.2 and Corollary 1.6.2 and is not included from my characterization of narrow bracketing. I justify such exclusion in two ways. From the theoretical point of view, computing the utility level of partial narrow bracketing necessitates computing the utility levels of expected utility and full narrow bracketing separately, and hence is arguably more complex and demanding. This contradicts with my interpretation that narrow bracketing is a simplifying heuristic for the multi-source risk. Empirically, Ellis and Freeman (2020) show that very few subjects are best described by non-trivial partial narrow bracketing with $\alpha \in (0,1)$. This suggests that incorporating such intermediate degrees of narrow bracketing might not significantly increase the model's explanatory power in practice.

1.6.1 Ranking of Rationality

In this section, I develop a ranking of rationality over models in Corollary 1.6.2 based on the readiness with which they yield violations of stochastic dominance.

I first define stochastic dominance on a subdomain of lotteries. A *choice environment* is a subset of lotteries $\mathcal{C} \subseteq \mathcal{P}$. A preference $\succsim$ is said to satisfy *stochastic dominance on* $\mathcal{C}$ if for any $P, Q \in \mathcal{C}$, $f[P] \succ_{FOSD} f[Q]$ implies that $P \succ Q$ and $f[P] = f[Q]$ implies that $P \sim Q$. Moreover, I say that the EU model satisfies stochastic dominance on $\mathcal{C}$ if for any $\succsim$ that can be represented by some V^{EU} in Corollary 1.6.2, $\succsim$ satisfies stochastic dominance on $\mathcal{C}$. Similar notions for other models including CN, ANB, ANB-CN and FNB can be defined.

For two choice environments $\mathcal{C}$ and $\mathcal{C}'$ with $\mathcal{C}' \subseteq \mathcal{C}$, if $\succsim$ satisfies stochastic dominance on $\mathcal{C}$, then by definition it also satisfies stochastic dominance on $\mathcal{C}'$. Hence I can define a complexity order $\trianglerighteq$ over choice environments such that $\mathcal{C}$ is more complex than $\mathcal{C}'$, i.e., $\mathcal{C} \trianglerighteq \mathcal{C}'$, if and only if $\mathcal{C}' \subsetneq \mathcal{C}$. My goal is to construct choice environments differing in the complexity levels in order to discriminate among the models characterized in Corollary 1.6.2.

The most complex choice environment is the set of all lotteries $\mathcal{P}$. Another natural choice environment is the riskless choice environment $\mathcal{C}_1 := Z$, which is the set of all degenerate lotteries. All five models in Corollary 1.6.2 satisfy stochastic dominance on $\mathcal{C}_1$. However, $\mathcal{C}_1$ is non-trivial in the sense that models featuring summation of expected utilities like $\hat{V}^{FNB}$ violate stochastic dominance on it. I define the following three choice environments whose complexity levels lie between $\mathcal{P}$ and $\mathcal{C}_1$:

- $\mathcal{C}_2 := \mathcal{L}^0(X) \times X$ is the choice environment where there is no risk in source 2;
- $\mathcal{C}_3 := \left\{P \in \mathcal{P} : P_{2|x} \in X,\ \forall x \in X\right\}$ is the choice environment where there is no conditional risk in source 2 give the income in source 1;

- $\mathcal{C}_3' := \hat{\mathcal{P}} = \mathcal{L}^0(X)^2$ is the choice environment where there is no correlation between income from two sources;

Denote $\mathcal{C}_4 := \mathcal{P}$. I have the following complexity ranking of choice environments:

$$\mathcal{C}_4 \trianglerighteq \mathcal{C}_3, \mathcal{C}_3' \trianglerighteq \mathcal{C}_2 \trianglerighteq \mathcal{C}_1,$$

where $\mathcal{C}_3$ and $\mathcal{C}_3'$ cannot by ranked by $\trianglerighteq$.

The next proposition shows that the above five complexity environments can be used to distinguish the five models in Corollary 1.6.2.

Corollary 1.6.2.

Proposition 1.6.3 *(i). FNB satisfies stochastic dominance on $\mathcal{C}_1$, but not on $\mathcal{C}_2$;*

(ii). ANB-CN satisfies stochastic dominance on $\mathcal{C}_2$, but not on $\mathcal{C}_3$ or $\mathcal{C}_3'$;

(iii). ANB satisfies stochastic dominance on $\mathcal{C}_3$, but not on $\mathcal{C}_3'$;

(iv). CN satisfies stochastic dominance on $\mathcal{C}_3'$, but not on $\mathcal{C}_3$;

(v). EU satisfies stochastic dominance on $\mathcal{C}_4$.

Model A is said to be *more rational* than model B if model A satisfies stochastic dominance on a more complex choice environment than model B. This rationality ranking reflects how readily a model yields violations of stochastic dominance. Proposition 1.6.3 reveals the rationality ranking of five models as depicted in Figure 1.2. In the graph, model A is more rational than model B if and only if there is a path from A to B. Hence, EU is more rational than CN and ANB, which are more rational than ANB-CN, while FNB is the least rational.

The rationality ranking in Figure 1.2 is consistent with the degree of behavioral heuristics employed in each model as shown in Figure 1.1. The more and stronger heuristics a model entails, the more it deviates from the expected utility benchmark, the more likely it is to yield violations of stochastic dominance and the more irrational

Figure 1.2: Rationality ranking of models based on violations of stochastic dominance.

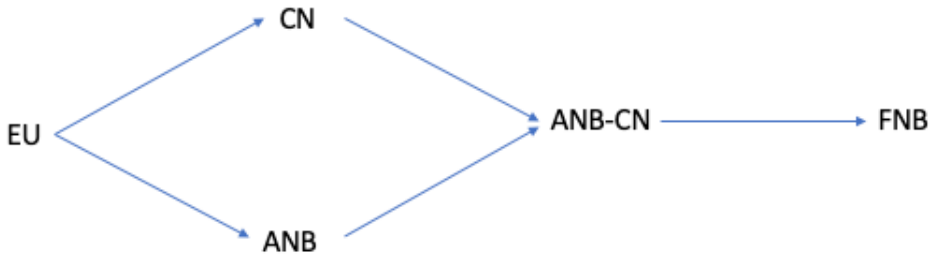

it is. This also explains why CN and ANB cannot be ranked in terms of rationality, because they adopt disjoint sets of behavioral heuristics.

1.6.2 Risk Aversion over Small Gambles

In this section, I apply narrow bracketing to problems with background risk. Consider a decision maker who is contemplating whether to purchase small-scale insurance for her new smart phone. Since the risk of a smart phone is negligible compared to the risk in the life-time wealth, a rational agent should take into account this background risk and be almost risk neutral regarding the risk of a smart phone. The popularity of small-scale insurance in practice reveals the prevalence of narrow bracketing when there is background risk.

Below I briefly show that narrow bracketing can explain risk aversion over small gambles. Consider the calibration result of Rabin (2000) which states that if an EU maximizer turns down 50-50 gambles of losing \$1000 or gaining \$1050 for all initial wealth levels, then she would always turn down 50-50 gambles of losing \$20,000 or gaining any sum. This serves as a critique on the expected utility theory since the DM has plausible risk aversion over small gambles, but is implausibly risk averse over large gambles.

Suppose that the DM's preference is represented by V^{FNB} with the utility index u given by equation (1.2) instead of V^{EU}. Then I can show that she will also reject the above small gambles, but she is willing to accept a 50-50 gamble of losing \$20,000

or gaining slightly more than \$80,000. Similar results can be derived if source 1 represents the background wealth and source 2 represents the outcome of the gamble, and the DM's preference is represented by an asymmetric narrow bracketing (with or without correlation neglect) utility function.

1.7 Applications: Time Preferences

In this section, I study the implications of narrow bracketing and correlation neglect in time preferences. Each outcome profile represents a consumption path in two periods $t = 1, 2$. A DM who finds it complex to evaluate intertemporal risk might resort to heuristics. For simplicity, I assume that the consumption space in each period is a compact interval $C := [\underline{c}, \bar{c}] \subset \mathbb{R}_+$ and focus on preferences that satisfy the following assumption.

Discounted Utility without Risk: There exists a continuous and strictly increasing function $u : C \to \mathbb{R}$ and $\beta \in (0, 1)$ such that for all $x_1, x_2, y_1, y_2 \in C$, $(x_1, x_2) \succsim (y_1, y_2)$ if and only if $u(x_1) + \beta u(x_2) \geq u(y_1) + \beta u(y_2)$.

This assumption follows from Dillenberger, Gottlieb, and Ortoleva (2020) and it postulates that when there is no risk, preferences can be represented by the summation of discounted utilities of consumptions in different periods. This holds for most models of time preferences in the literature.

Corollary 1.7.1 *Suppose that $\succsim$ satisfies the axioms stated in Theorem 1.4.2 and can be represented by one of V^{EU}, V^{ANB}, V^{CN}, V^{ANB-CN} or V^{FNB}. Then $\succsim$ satisfies Discounted Utility without Risk if and only if the index $w : C^2 \to \mathbb{R}$ in each utility function can be rewritten as*

$$w(x_1, x_2) = \phi(u(x_1) + \beta u(x_2)), \forall (x_1, x_2) \in C^2,$$

where $\phi : [(1+\beta)u(\underline{c}), (1+\beta)u(\overline{c})] \to \mathbb{R}$ *is continuous and strictly monotonic.*

Below I show that different utility functions in Corollary 1.7.1 correspond to different commonly used time preferences models in the literature.

Example 1.7.2 (EU). V^{EU} takes the form of

$$V^{EU}(P) = \sum_{x_1,x_2} \phi(u(x_1) + \beta u(x_2))P(x_1, x_2),$$

which is the Kihlstrom-Mirman (KM) representation studied in Kihlstrom and Mirman (1974) and Dillenberger, Gottlieb, and Ortoleva (2020). If ϕ is affine, then $\succsim$ can be represented by $\sum_{x_1} u(x_1)P_1(x_1) + \beta \sum_{x_2} u(x_2)P_2(x_2)$, the standard Expected Discounted Utility (EDU) model.[2]

Example 1.7.3 (FNB). Suppose $v_1 = v_2 := v$, that is, narrow preferences in both sources are identical, then V^{FNB} takes the form of

$$V^{FNB}(P) = u(CE_v(P_1)) + \beta u(CE_v(P_2)).$$

The DM first narrowly evaluates the marginal risk in each period and then aggregates the certainty equivalents using discounted utility. Note that I can normalize ϕ to be the identity function since w is unique up to a monotonic transformation in V^{FNB}. This functional form is the adaption of the Dynamic Ordinal Certainty Equivalent (DOCE) model (Selden, 1978; Selden and Stux, 1978, Kubler, Selden, and Wei, 2020b) to my setup. However, it emerges from narrow bracketing in my model, instead of normative properties of time preferences.

Example 1.7.4 (ANB). V^{ANB} takes the form of

$$V^{ANB}(P) = \sum_{x_1} \phi(u(x_1) + \beta u(CE_v(P_{2|x}))P_1(x_1). \tag{1.3}$$

[2]EDU can also be derived from FNB in Example 1.7.3 by assuming $v = u$, and from ANB in Example 1.7.4 by assuming $\phi(x) = x$ for all x and $v = u$.

If I fix the consumption levels for periods 0 and $t > 2$, then this functional form corresponds to a majority of recursive preferences built upon Kreps and Porteus (1978). The DM evaluates risk in different periods recursively. Given the current consumption, she first reduces the conditional future risk to its certainty equivalent; then she evaluates this certainty equivalent and the current consumption using discounted utility; finally she moves one period ahead and repeats the above process. Like DOCE in Example 3, such features also emerge from narrow bracketing, instead of a preference for timing of resolution of uncertainty as in recursive preferences. In Section 1.7.1, I will illustrate this point in detail by comparing ANB in (1.3) with the Epstein-Zin (EZ) preferences in Epstein and Zin (1989) and Weil (1990).

1.7.1 Epstein-Zin Type Preferences

In this section, I focus on the ANB utility function in (1.3) and show that it incorporates EZ type preferences. That is, an EZ type utility function emerges because of narrow bracketing. Moreover, I will argue that this new foundation of EZ type preferences might be more suitable for some applications than the original foundation in Epstein and Zin (1989) since it can address the critique in Epstein, Farhi, and Strzalecki (2014).

To facilitate the comparison of ANB and EZ preferences, I assume that the consumption in the first period is deterministic and focus on the restriction of $\succsim$ on $C \times \mathcal{L}^0(C)$.[3] Consider the following specification of ANB

$$V^{ANB}(c_1, p) = \frac{1}{\alpha}\Big\{c_1^{\rho} + \beta[\mathbb{E}_p(c_2^{\alpha})]^{\rho/\alpha}\Big\}^{\alpha/\rho},\ \forall\ (c_1, q) \in C \times \mathcal{L}^0(C). \tag{1.4}$$

where $0 \neq \rho < 1, 0 \neq \alpha < 1$ and $0 < \beta < 1$. Note that $\gamma := 1 - \alpha \in \mathbb{R}_+\backslash\{1\}$ is the coefficient of relative risk aversion (RRA) and $\psi = \frac{1}{1-\rho} \in \mathbb{R}_+\backslash\{1\}$ is the coefficient

[3]In the appendix, I briefly discuss how to extend the ANB model to an infinite horizon setting.

of elasticity of intertemporal substitution (EIS). (1.4) can be derived from (1.3) by assuming $u(x) = \frac{x^\rho}{\rho}$, $v(x) = \frac{x^\alpha}{\alpha}$ and $\phi(x) = v \circ u^{-1}(x)$. When $\alpha = \rho$, i.e., RRA equals the reciprocal of EIS, (1.4) reduces to EDU.

Since $f(x) = \frac{x^\alpha}{\alpha}$ is strictly monotone, f^{-1} is well-defined and strictly monotone. I impose the monotonic transformation f^{-1} on V^{ANB} without changing the preference on $C \times \mathcal{L}^0(C)$:

$$\hat{V}^{ANB}(c_1, p) = \left\{c_1^\rho + \beta[\mathbb{E}_p(c_2^\alpha)]^{\rho/\alpha}\right\}^{1/\rho}, \ \forall \ (c_1, p) \in C \times \mathcal{L}^0(C). \tag{1.5}$$

By comparison, the standard CRRA-CES EZ preference entails the following recursive formulation of continuation utility U_t:

$$U_t = \left\{c_t^\rho + \beta\left[\mathbb{E}_t\left(U_{t+1}^\alpha\right)\right]^{\rho/\alpha}\right\}^{1/\rho}. \tag{1.6}$$

Note that the ANB utility function (1.5) is quite similar to the recursive EZ utility function (1.6). However, I cannot simply apply (1.6) to $C \times \mathcal{L}^0(C)$ since EZ is built upon a richer space than the set of lotteries over consumption paths. The richer space involves temporal resolution of uncertainty. Following Kreps and Porteus (1978), I call it the set of *temporal lotteries* and denote by d a generic temporal lottery. The difference between a temporal lottery and a lottery can be illustrated by the following simple example.

Figure 1.3: Example of temporal lotteries inducing the same lottery.

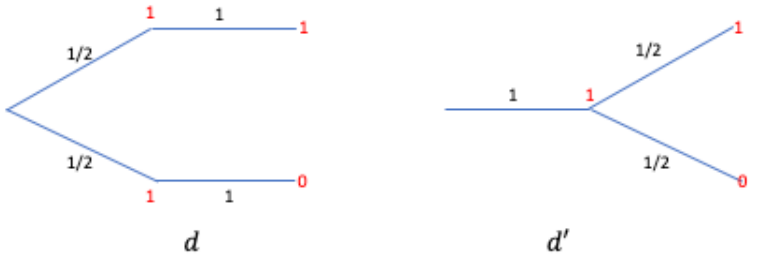

Consider the two temporal lotteries d and d' depicted in Figure 1.3. Both d and d' deliver consumption 1 for sure at date 1 and have probability 1/2 to deliver either consumption 1 or consumption 0 at date 2; that is, d and d' induce the same lottery over consumption paths. However, they differ in timing of resolution of risk. In temporal lottery d, the consumer knows her future consumption at date 1, while in temporal lottery d', the risk about consumption at date 2 is only resolved at date 2.

EZ captures the consumer's attitude towards the above difference in timing of resolution of risk. Formally, for any lottery $(c_1, p) \in C \times \mathcal{L}^0(C)$, I denote by d^E the temporal lottery that induces lottery (c_1, p) and has early resolution of risk, like d in Figure 1.3. Similarly, I denote by d^L the temporal lottery that induces lottery (c_1, p) and has late resolution of risk, like d' in Figure 1.3. According to the EZ functional form in (1.6), the EZ utility of d^E and d^L are given by

$$V^{EZ}(d^E) = \Big\{\mathbb{E}_p\Big[(c_1^\rho + \beta c_2^\rho)^{\alpha/\rho}\Big]\Big\}^{1/\alpha},$$
$$V^{EZ}(d^L) = \Big\{c_1^\rho + \beta[\mathbb{E}_p(c_2^\alpha)]^{\rho/\alpha}\Big\}^{1/\rho}.$$

One can easily show that if (c_1, p) is induced by d^E and d^L, then $V^{EZ}(d^L) = \hat{V}^{ANB}(c_1, p)$ and $V^{EZ}(d^E)$ agrees with a monotone transformation of $V^{EU}(c, p)$ with the same preference parameters, i.e., $V^{EU}(c, p) = \frac{1}{\alpha}\mathbb{E}_p\Big[(c_1^\rho + \beta c_2^\rho)^{\alpha/\rho}\Big]$. Hence, the ANB utility function $\hat{V}^{ANB}$ can be interpreted as the restriction of the EZ utility function V^{EZ} on temporal lotteries where there is no early resolution of uncertainty.

It is well-known that when $\rho > \alpha$, that is, when RRA $>$ 1/EIS, a consumer with the EZ utility function exhibits a preference for early resolution of risk and strictly prefers d to d'. Symmetrically, she has a preference for late resolution of risk when $\rho < \alpha$. As a result, the consumer is indifferent to timing of risk resolution and attaches the same value to d and d' if and only if $\alpha = \rho$, in which case EZ reduces to

EDU. Hence the EZ utility function separates the time and risk preferences only if it entails a strict preference for either early or late resolution of risk.

In contrast, although EZ type preferences can also take the form of ANB as in (1.5), it emerges from narrow bracketing, instead of a non-trivial attitude towards timing of risk resolution. Actually, since d and d' induce the same lottery over consumption paths, a consumer with ANB utility function (1.5) is naturally indifferent between d and d' and hence exhibits indifference to timing of resolution of risk.

My alternative foundation of EZ type preferences based on narrow bracketing has the following implications. Theoretically, it shows that one can simultaneously allow for indifference to temporal resolution of risk and separation of time and risk preferences, which is impossible in the original foundation in Epstein and Zin (1991) based on the framework of Kreps and Porteus (1978). It reveals a novel and surprising connection between the EZ type utility function, the most commonly used alternative to EDU, and narrow bracketing, the behavioral heuristic to simplify the evaluation of multi-source risk. Finally, in order to perform a thorough test of the behavioral implications of an EZ-type utility function, my foundation only requires the econometrician to observe the DM's choices in the natural domain of lotteries over consumption paths, while the original foundation entails choices in the more complex domain of temporal lotteries.

The difference of the two utility functions is also empirically relevant and important. Thanks to the separation of time and risk preferences and the convenient functional form, EZ has been widely applied in macroeconomics and finance. For instance, the long-run risks model of Bansal and Yaron (2004) considers a representative agent with CRRA-CES EZ preferences and delivers a unified explanation for several long-standing puzzles of asset markets, including the equity premium puzzle, the excessive asset price volatility, and the large cross-sectional differences in average returns across equity portfolios. Their main empirical results are based on parameters

with a RRA of 10 and an IES of 1.5, but have ignored the quantitative implications of the strong preference for early resolution of risk. Through introspection, Epstein, Farhi, and Strzalecki (2014) show that the preference parameters and calibrated endowment process imply that the representative agent in the long-run risks model is willing to give up 25 or 30 percent of her lifetime consumption in order to have all risk about future consumption resolved in the next period. This timing premium is unrealistically high since the risk is about future consumption instead of future income or asset returns, and hence such information has no apparent instrumental value. In other words, the representative agent has no need to reoptimize her contingent consumption plans given early resolution of risk. Meissner and Pfeiffer (2022) show that on average subjects give up 5% of their total consumption to resolve all uncertainty immediately in an experiment.[4] Epstein, Farhi, and Strzalecki (2014) also show that the high timing premium is robust to other models using EZ preferences (such as the rare disasters model in Barro (2009)), different preference parameter values and more general risk preferences. In contrast, in the EZ type preferences based on ANB, the representative agent attaches no value to the non-instrumental information and the timing premium is always zero.

I conclude the section with a final comment. Note that ANB only specifies the preference on lotteries over consumption paths. Although the most natural way to extend ANB to temporal lotteries is to assume indifference to timing of resolution of risk as I did in this chapter, one can also extend ANB by assuming a preference for early or late resolution of uncertainty, which I leave for future research. The main message of my alternative foundation of EZ type preferences is that the attitude towards temporal resolution of uncertainty can be isolated from the separation of time and risk preferences.

[4]Meissner and Pfeiffer (2022) also show a negative correlation between predicted and elicited timing premia in EZ, which questions the structural connection between the preference for early resolution of uncertainty and the preference parameters in EZ.

1.7.2 Effects of Heuristics on Optimal Saving

In this section, I illustrate how narrow bracketing and correlation neglect affect the saving behavior of a consumer who faces income uncertainty and deterministic interest rates through a parsimonious two-period model. One can interpret my analysis as comparative statics of optimal saving for simplifying heuristics, similar to the comparative statics for future income risk or risk attitude in the literature.

Following Selden (1978) and many papers on precautionary savings including Drèze and Modigliani (1972) and Kimball (1990), I start with the standard "certain × uncertain" setup, where there is only one saving decision (in the first period) and income is uncertain in a single period (in the second period). Concretely, at date 1, the consumer receives income $y_1 > 0$ which is certain and can be allocated between consumption and saving. At date 2, state $i \in \{1, \dots, N\}$ is realized with probability $\pi_i \in (0, 1)$ and $\sum_{i=1}^{N} \pi_i = 1$, which determines the level of date 2 income $y_{2,i} > 0$ with $y_{2,i} \neq y_{2,j}$ for $i \neq j$. The gross return on savings $R > 0$ is independent of the state at date 2. The consumer needs to choose the saving s at date 1 before observing the state of the world. I allow s to be negative, in which case the consumer borrows at date 1 with interest rate R, given that her future income in any realized state can cover the debt. Denote c_1 as the consumption level at date 1 and $c_{2,i}$ as the consumption level at date 2 when state i occurs. The budget constraints are given by

$$\begin{aligned} y_1 - s = c_1 &\geq 0, \\ y_{2,i} + Rs = c_{2,i} &\geq 0 \text{ for all } i = 1, \dots, N. \end{aligned}$$

Each choice of saving s induces a lottery over 2-period consumption paths P such that P_1 is degenerate at c_1 and $P_2(c_{2,i}) = \pi_i$ for each $i = 1, \dots, N$. Hence, it suffices to specify the consumer's preference over $C \times \mathcal{L}^0(C) \subset \hat{\mathcal{P}}$. Given the same preference parameters, among the different utility functions characterized in Corollary 1.7.1,

V^{EU} and V^{CN} coincide on $C \times \mathcal{L}^0(C)$, and V^{ANB}, V^{ANB-CN} and V^{FNB} coincide on $C \times \mathcal{L}^0(C)$. I denote the former by V^{EU} and the later by V^{NB}. [5] This observation suggests that correlation neglect has no bite in this optimal saving problem and I can focus on the effects of narrow bracketing.

For computational simplicity, I further assume that V^{NB} admits a CRRA-CES version with unitary EIS, where I apply (1.3) to $C \times \mathcal{L}^0(C)$ with $u(x) = \log x$, $v(x) = \frac{x^\alpha}{\alpha}$ and $\phi(x) = v \circ u^{-1}(x)$ with $\alpha < 1$ and $\alpha \neq 0$:

$$V^{NB}(c_1, q) = \frac{1}{\alpha} c_1^\alpha \cdot \Big[\sum_{c_2} c_2^\alpha q(c_2) \Big]^\beta, \ \forall \ (c_1, q) \in C \times \mathcal{L}^0(C).$$

Using the same preference parameters, the expected utility benchmark is V^{EU} where

$$V^{EU}(c_1, q) = \frac{1}{\alpha} c_1^\alpha \cdot \sum_{c_2} c_2^{\beta\alpha} q(c_2), \ \forall \ (c_1, q) \in C \times \mathcal{L}^0(C).$$

Denote by s^{EU} and s^{NB} the optimal levels of savings if the consumer's preference is represented by V^{EU} and V^{NB} respectively. The difference between s^{EU} and s^{NB} reflects the impact of narrow bracketing on the optimal saving behavior, isolated from the preference parameters and correlation neglect.

Lemma 1.7.5 (Saving with narrow bracketing) *Consider the "certain × uncertain" saving problem above. Then $s^{NB} > s^{EU}$ if $\alpha < 0$ and $s^{NB} < s^{EU}$ if $\alpha > 0$.*

The intuition behind Lemma 1.7.5 is straightforward. With the same preference parameters, the effective coefficient of RRA for consumption risk in period 2 is $1 - \beta\alpha$ in the EU benchmark, and $1 - \alpha$ in the model with narrow bracketing. Since the discount factor $\beta \in (0, 1)$, when $\alpha < 0 = \rho$, I have $1 - \alpha > 1 - \beta\alpha$, which suggests that narrow bracketing makes the consumer effectively more risk averse regarding the

[5] The observation that ANB and FNB, and hence the models in Kreps and Porteus (1978) and Selden (1978), agree in the "certain × uncertain" setup also appears in Bommier, Chassagnon, and Le Grand (2012) and Kubler, Selden, and Wei (2020a).

consumption risk in period 2. This intuition remains valid for the general CRRA-CES model in (1.4) with $\rho > \alpha$. Although it is well-known that the relationship between the optimal saving level and the risk aversion parameter α is non-monotonic in the EZ model,[6] I show that the increased effective risk aversion due to narrow bracketing does lead to higher saving in my setup. Since the preference over deterministic consumption paths are the same in V^{NB} and V^{EU}, my result also implies that narrow bracketing increases precautionary saving when $\alpha < 0$.

Recall that RRA is $1 - \alpha$ and EIS is $\frac{1}{1-\rho}$. For $\rho = 0$, RRA $>$ 1/EIS if and only if $\alpha < 0$. In order to explain the equity premium puzzle and other puzzles in the financial market using the EZ preference, the empirical asset pricing literature, such as the long-run risks model (Bansal and Yaron, 2004) and the rare disasters model (Barro, 2009), relies on parameters satisfying RRA $>$ 1/EIS. My Lemma 1.7.5 provides some intuition why this combination of parameters can work based on the idea that an EZ type utility function emerges because of narrow bracketing. When RRA $>$ 1/EIS, a consumer with narrow bracketing is effectively more risk averse regarding future consumption risk compared to a consumer with EU and the same RRA parameter $1 - \alpha$. As a result, the EZ type utility function can explain the excessive equity premium with the RRA parameter $1 - \alpha$ lying in a reasonable range.

In order to study the effect of the other heuristic, correlation neglect, on the consumer's optimal saving behavior, I need to modify the previous setup to allow for uncertainty over consumption in the first period. I assume that income is uncertain in the second period, but the consumer knows her future income in the first period and can tailor her saving decision to such information. Although the consumer is indifferent to temporal resolution of uncertainty and does not attach any value to non-instrumental information, she does value such instrumental information about the future income since it can improve the quality of her saving decisions in the

[6]See Kimball and Weil (2009) and Bommier, Chassagnon, and Le Grand (2012). Bommier, Kochov, and Le Grand (2017) illustrates this non-monotonic relationship in a simple example.

first period. For simplicity, I further assume that the income in the second period is either $y_{2,h}$ or $y_{2,l}$, each with probability 1/2 and $y_{2,h} > y_{2,l} > 0$. I will refer to this as the "uncertain $\times$ conditionally certain" saving problem, since there is no income uncertainty date 2 given the information in date 1.

In this "uncertain $\times$ conditionally certain" setup, the consumer needs to choose the saving level s_i at date 1 conditional on the state $i = h, l$. Denote by $c_{1,i}$ $c_{2,i}$ the consumption levels at date 1 and date 2 respectively when state i occurs. Then the budget constraints are

$$y_1 - s_i = c_{1,i} \geq 0 \text{ for } i = l, h,$$
$$y_{2,i} + Rs_i = c_{2,i} \geq 0 \text{ for } i = l, h.$$

Each choice of saving levels s_l and s_h induces a lottery over consumption paths P where $P(c_{1,l}, c_{2,l}) = P(c_{1,h}, c_{2,h}) = 1/2$. If $s_h \neq s_l$, then $c_{1,h} \neq c_{1,l}$ and there is risk about consumption at date 1 but no conditional risk about consumption at date 2. If instead $s_h = s_l$, then $c_{1,h} = c_{1,l}$ and the consumption at date 1 is certain and there is risk about consumption at date 2.

I assume that the consumer either has expected utility or adopts the correlation neglect heuristic, that is, her preference $\succsim$ is represented by one of the following two utility functions

$$V^{EU}(P) = \frac{1}{\alpha} \sum_{c_1,c_2} c_1^{\alpha} \cdot c_2^{\beta\alpha} P(c_1, c_2),$$
$$V^{CN}(P) = \frac{1}{\alpha} \sum_{c_1,c_2} c_1^{\alpha} \cdot c_2^{\beta\alpha} P_1(c_1) P_2(c_2).$$

Denote by s_i^{EU} and s_i^{CN}, $i = l, h$, the optimal levels of savings in two states for the utility function V^{EU} and V^{CN} correspondingly. The different between s_i^{CN} and

s_i^{EU} reflects the impact of correlation neglect on the optimal saving behavior in state $i = l, h$, isolated from the preference parameters and narrow bracketing.

Lemma 1.7.6 (Saving with correlation neglect) *Consider the "uncertain $\times$ conditionally certain" saving problem described above. The following results hold:*

(i). If $\alpha < 0$, then $s_h^{EU} < s_h^{CN} < s_l^{CN} < s_l^{EU}$;

(ii). If $\alpha > 0$, then $s_h^{CN} < s_h^{EU} < s_l^{EU} < s_l^{CN}$.

The effects of correlation neglect on the optimal saving behavior depends on whether the consumer is correlation averse or correlation seeking in the EU benchmark. Following Bommier (2007), I say the consumer is *correlation averse* if and only if for all $x_1, x_2, y_1, y_2 \in C$ with $x_1 < x_2$ and $y_1 < y_2$, the lottery $\frac{1}{2}(x_1, y_2) + \frac{1}{2}(x_2, y_1)$ is preferred to the lottery $\frac{1}{2}(x_1, y_1) + \frac{1}{2}(x_2, y_2)$. In other words, a correlation averse consumer prefers consumption at different dates to be negatively correlated rather than positive correlated. For V^{EU} considered in this section, the consumer is correlation averse if and only if $\alpha < 0$. Compared to a consumer with correlation neglect, a correlation averse consumer adjusts the current saving levels to compensate for the difference in future income. She saves more when the future income is low and less when the future income is high. A symmetric argument holds for the comparison between correlation seeking and correlation neglect.

1.8 Conclusion

In this chapter, we study narrow bracketing and correlation neglect as two heuristics that decision makers adopt to simplify the evaluation of risk from multiple sources. By relaxing the independence axiom of the expected utility model, we develop a class of models that (i) allow for narrow bracketing or correlation neglect or both and (ii) allow for different degrees of deviations from the benchmark without narrow bracketing. We interpret the different sources as different streams of income to explain

experimental evidence on violations of first order stochastic dominance. We propose a rationality ranking of our models based on how likely they are to yield violations of stochastic dominance. Then, we interpret one source as background risk and show that narrow bracketing can explain risk aversion over small gambles. Finally, we interpret the different sources as consumption in different periods to explore the implications of narrow bracketing and correlation neglect on the optimal saving behavior of a consumer. We also show that an Epstein-Zin type utility function can emerge because of narrow bracketing. This new utility function achieves separation between time and risk preferences without inducing an implausibly high timing premium as observed by Epstein, Farhi, and Strzalecki (2014). My results show that both the separation of time and risk preferences and violations of stochastic dominance can result from narrow bracketing and both can be captured by relaxing the independence axiom.

Our analysis suggests several promising directions for future research. In the appendix, we briefly discuss how to extend the asymmetric narrow bracketing model in an infinite horizon setting. We leave for future research the axiomatization of narrow bracketing with more than two sources of risk. In another ongoing project, we consider general models of correlation misperception by modeling the correlation structure among risk in difference sources as copula. Moreover, we maintain the independence axiom within each source and focus on narrow bracketing and correlation neglect as simplifying heuristics for inter-source risk. One extension is to deviate from expected utility within each source and incorporate other commonly used behavioral factors under one-source risk, like cumulative prospect theory and certainty effect.

Chapter 2

Maxmin Implementation

2.1 Introduction

[1]Starting from Knight (1921), Keynes (1921) and Ellsberg (1961), it has been argued that ambiguity aversion plays an important role in individual decision making. Experimental evidence suggests that decision makers might avoid choosing prospects containing ambiguous components.[2] To account for decision makers' ambiguity aversion, various theories have been proposed, among which the maxmin expected utility (MEU) theory is one of the most notable ones, and has been applied in studying economic problems in different fields.[3] In this chapter, we explore the implementation problem of choice correspondences with MEU agents and demonstrate how to exploit the agents' ambiguity aversion.

Consider a mechanism designer (MD) and one or multiple agents. Without confusion, we will refer to the MD as *she* and each agent as *he*. The MD desires to implement a choice correspondence which maps each type profile of the agents to a nonempty set of outcomes. One interpretation is that, given a type profile of the agents, the MD is indifferent among all outcomes in the mapped set, and thus does not care about which one is chosen from the set. The multiplicity of the MD's desired outcomes reflects the potentially coarse nature of the MD's objective. For instance, a government who plans to build railway roads to connect cities A, B and C might be indifferent between whether to connect A, B and B, C or to connect A, C and A, B if both are socially efficient. Implicitly, we assume that the MD has a preference over the outcomes, which is possibly incomplete and determined by the type profile of the agents. The set of outcomes mapped by the choice correspondence contains all undominated outcomes given the MD's preference. For the main part of this

[1]This chapter is based on the paper "Maxmin Implementation" joint with Rui Tang. It has been published in *Journal of Economic Theory*, 194: 105250, 2021.

[2]See, for instance, Halevy (2007) and Chew et al. (2017).

[3]The MEU model is introduced and axiomatized by Gilboa and Schmeidler (1989). For its applications, see, for instance, Epstein and Schneider (2008a), Castro and Yannelis (2018), etc.

chapter, we abstract away the MD's preference and directly work with a given choice correspondence.

We consider the quasi-linear environment where the MD can incentivize the agents through monetary transfers. A *single-mechanism* consists of an allocation rule and a transfer scheme. The allocation rule maps each type profile reported by the agents to a distribution over outcomes. The transfer scheme specifies the transfers paid by each agent based on the reported type profile. A *multi-mechanism* consists of a nonempty set of single-mechanisms. By using a multi-mechanism, the MD commits to one of the single-mechanisms constituting the multi-mechanism without revealing it. As a result, being uncertain about which specific single-mechanism is adopted, a MEU agent evaluates his payoff according to the infimum of all possible payoffs he could get from the set of single-mechanisms.

Prior to introducing the definition of implementability of a choice correspondence, we point out that certain extent of ambiguity is used in real-world mechanisms. One example is that Google's auction of advertisement space on its search result pages. Google states that the bid, advertisement quality and advertisement format will jointly determine the appearance and positioning of an advertisement, but the exact rules are not publicly revealed.[4] Another prominent example is the commonly agreed rule *Contra proferentem*, which states that ambiguity in a written instrument should be construed most strongly against the party responsible for the choice of language. As a result, as long as the MD can validate that she cannot take advantage of the ambiguity of the mechanism, the adoption of ambiguity is justified.[5] In our framework, a MD can validate her use of a multi-mechanism if every single-mechanism

[4]See, for example, `https://support.google.com/google-ads/answer/6366577?hl=en`.

[5]For instance, in 2013, a firm in Scotland named HFD Construction Limited accused the Aberdeen City Council of ambiguity in the tender documents during the procurement process. HFD's bid was unsuccessful and it had interpreted the requirements differently from the winner. The court refused HFD's petition and argued that different interpretations were reasonable given that the Council was seeking and encouraging innovative proposals for the development of the local economy. For details of the case, please refer to `https://www.scotcourts.gov.uk/search-judgments/judgment?id=0a5286a6-8980-69d2-b500-ff0000d74aa7`.

of the multi-mechanism implements the MD's desired outcomes. This is indeed one of the requirements in our definition of implementability.

We say that a choice correspondence is implementable if there exists a multi-mechanism such that (1) truth-telling constitutes an equilibrium if each agent evaluates his payoff according to the worst possible payoff from the set of single-mechanisms, (2) the equilibrium payoff of each agent is not negative infinity, and (3) each single-mechanism maps each type profile of the agents to a distribution whose support is in the set of MD's desired outcomes, i.e., the chosen outcome by each single-mechanism is in the set mapped by the choice correspondence almost surely. While condition (2) is a technical requirement, conditions (1) and (3) require that the MD has no incentive to cherry-pick a particular single-mechanism from the ones constituting the multi-mechanism and that each agent has incentive to report truthfully given that the others are doing so. With MEU agents, a MD can use a multi-mechanism to implement a choice correspondence which cannot be implemented by a single-mechanism. We illustrate this point through the following simple example.

Example 2.1.1 Consider a government who wants to delegate the construction of a bridge to one of two firms indexed by $a \in \{1, 2\}$. Each firm a's type is given by a two-dimensional vector $\theta^a = (q^a, c^a) \in \{1, 0\} \times \{1, 0\}$. q^a represents the firm's construction quality where $q^a = 1$ means high quality and $q^a = 0$ means low quality. c^a measures the firm's construction cost, which is either 1 or 0. The types of the two firms are private information and have independent and identical distribution P. The distribution P satisfies that $P(1, 0) = 0.4$ and $P(1, 1) = P(0, 1) = P(0, 0) = 0.2$. Each firm's payoff is the transfer from the government minus the firm's cost of construction if he wins and 0 otherwise. The government does not care about the transfers and has a lexicographic preference: she prefers firms with higher construction quality, and in case of a tie of construction quality, she chooses the less costly one. Let $\succ^G$ denote

the preference of the government, and we have $(1,0) \succ_G (1,1) \succ_G (0,0) \succ_G (0,1)$. The choice correspondence of the government is thus given by

$$F(\theta^1, \theta^2) = \Big\{ a \in \{1,2\} : q^a > q^b \text{ or } q^a = q^b, c^a \leq c^b, b \in \{1,2\}, b \neq a \Big\}.$$

We first argue that the choice correspondence F cannot be implemented by a single-mechanism. Suppose to the contrary that there exists such a single-mechanism. Obviously, the single-mechanism must select the government's preferred firm if the reported types of the two firms differ. Thus, only the tie-breaking rule of the single-mechanism needs to be specified. For any $\theta \in \{0,1\} \times \{0,1\}$, let $b_\theta \in [0,1]$ denote the probability of firm 1 being chosen when both firms report type θ. By standard arguments, each firm's interim probability of winning should be weakly decreasing with respect to the firm's construction cost. Thus, firm 1's interim winning probability of reporting type $(0,0)$ should not be smaller than that of reporting type $(1,1)$, i.e., $0.2 + 0.2b_{(0,0)} \geq 0.4 + 0.2b_{(1,1)}$. It implies that $b_{(0,0)} = 1$ and $b_{(1,1)} = 0$. By checking the same monotonicity constraint for firm 2, we have a contradiction. As a result, no single-mechanism can implement F.

In contrast, there exists a multi-mechanism that implements F if both firms are MEU maximizers. The government commits to a multi-mechanism containing two single-mechanisms. The two allocation rules only differ in tie-breaking rules in the sense that they both select the government's preferred firm and in case of a tie, the first one always selects firm 1 and the second one always selects firm 2. Expected transfers paid from the government to the two firms are shown in the following tables.

Table 2.1: Expected payments to the firms in the first single-mechanism

Type	$(1,0)$	$(1,1)$	$(0,0)$	$(0,1)$
Firm 1	0.4	0.6	0.4	0.2
Firm 2	0.4	0.4	0.4	0

Table 2.2: Expected payments to the firms in the second single-mechanism

Type	$(1,0)$	$(1,1)$	$(0,0)$	$(0,1)$
Firm 1	0.4	0.4	0.4	0
Firm 2	0.4	0.6	0.4	0.2

Now we check the incentive compatibility conditions between types $(1,1)$ and $(0,0)$ for firm 1 (and firm 2, due to the symmetry of the multi-mechanism). First, suppose that firm 1's type is $(1,1)$. Firm 1 gets payoff 0 in both single-mechanisms by reporting the true type $(1,1)$. To see this, in the first single-mechanism, firm 1's interim probability of winning is 0.6 since he wins in case of a tie. Hence, his expected cost is 0.6, which equals the expected payment from the government. Similarly, both the expected earning and the expected cost of firm 1 are 0.4 in the second mechanism by truthfully reporting the type $(1,1)$. If firm 1 reports his type as $(0,0)$, he gets payoff 0 under the first single-mechanism and payoff 0.2 under the second single-mechanism. Since firm 1 is a MEU maximizer, he evaluates his payoff by misreporting $(0,0)$ as 0 and thus has no strict incentive to do so. Second, suppose that firm 1 has type $(0,0)$. He receives payoff 0.4 under each single-mechanism by truth-telling. By misreporting $(1,1)$, he gets payoff 0.6 under the first single-mechanism and 0.4 under the second single-mechanism. Again, ambiguity aversion guarantees that firm 1 is indifferent between truth-telling and misreporting. Note that different types' incentive compatibility is guaranteed by different single-mechanisms. As a result, the monotonicity constraint for the interim winning probabilities is no longer necessary for each single-mechanism. One can simply check that other incentive compatibility conditions also hold.

The main result of our chapter characterizes implementable choice correspondences. For ease of illustration, we focus on the case with one agent in Section 2.3 and extend the result to the multi-agent case in Section 2.4. Our result general-

izes an existing result known for the implementability condition of choice functions: Rockafellar (1970) and Rochet (1987) show that a choice function f is implementable if and only if it satisfies *cyclical monotonicity*, i.e., for any finite sequence of types $\{\theta_1, ..., \theta_n\}$ with $n \geq 1$ and $\theta_{n+1} := \theta_1$,

$$\sum_{k=1}^{n} [u(\theta_{k+1}, f(\theta_{k+1})) - u(\theta_k, f(\theta_{k+1}))] \geq 0,$$

where $u(\theta, x)$ is the utility of a type θ agent when outcome x is chosen. We extend cyclical monotonicity condition to choice correspondences. We say that a choice correspondenc F satisfies *cyclical monotonicity* if for any finite sequence of types $\{\theta_1, ..., \theta_n\}$ with $n \geq 1$ and $\theta_{n+1} := \theta_1$,

$$\sum_{k=1}^{n} \sup_{x \in F(\theta_{k+1})} [u(\theta_{k+1}, x) - u(\theta_k, x)] \geq 0.$$

Under some boundedness condition, we prove that cyclical monotonicity is sufficient and necessary for a choice correspondence to be implementable. The boundedness condition is trivially satisfied when F reduces to a choice function. Thus, our result generalizes Rockafellar (1970) and Rochet (1987).

One feature of our result is that the set of implementable choice correspondences remains the same whether or not randomized reports are allowed. When an agent is a MEU maximizer, a randomized report might yield a strictly higher payoff for him than any deterministic report. Thus, restricting to pure strategies is with loss of generality for assessing incentive compatibility of a particular mechanism. In the literature, randomized reports are either directly excluded (Di Tillio et al., 2017) or proved to affect the implementability results (Bose and Renou, 2014).[6] In contrast, we show that

[6]Bose and Renou (2014) focus on deterministic reports in the paper and extend the characterization results by allowing for randomized reports in the Supplemental Material. They point out that the choice function in the introductory example, which is implementable when only deterministic reports are allowed, is no longer implementable when randomized reports are considered.

no matter whether randomized reports are allowed or not, cyclical monotonicity is equivalent to implementability of a choice correspondence. In Section 2.7.2, we impose an additional restriction that each single-mechanism yields the same expected total transfer. When there are at least two agents, our characterization results remain valid no matter whether randomized reports are allowed or not. However, for the single-agent case, it will make a difference whether or not the agent can choose a mixed strategy.

We explore our characterization result in more specialized settings including supermodular environments and one-dimensional type spaces in Section 2.5. We then apply the result to the allocation problem of a single good in Section 2.6, where a MD wants to allocate the good to her most desired agent. We allow the MD's preference to be incomplete, and prove that the implementability of the MD's optimal choice correspondence is equivalent to a quasi-monotonicity condition. The condition says that given the MD's preference, the highest possible interim winning probability of an agent when he has a higher private value (e.g., a lower construction cost) should be no smaller than his lowest possible interim winning probability when he has a lower private value (e.g., a higher construction cost). In the multi-mechanism, different interim winning probabilities of a fixed agent are given by different tie-breaking rules used by the MD. As is shown by Example 2.1.1, by using multi-mechanisms rather than single-mechanisms, the set of implementable choice correspondences can be strictly expanded.

Related Literature. Originating from the seminal work of Bergemann and Morris (2005) and Chung and Ely (2007), robust mechanism design has been widely studied. One motivation of this branch of literature is that the MD is ambiguous towards some critical components of the agents, including higher order beliefs (Chen and Li, 2018), available actions (Carroll, 2015), available information (Du, 2018; Brooks and Du,

2020), correlation of value distributions of multiple goods (Carroll, 2017), etc. As a result, the MD adopts the maxmin criterion to evaluate a mechanism.

In contrast, we consider a MD who intentionally creates ambiguity for MEU agents to implement her desired outcomes. Thus, our approach connects to the literature studying ambiguity averse agents, including Bose et al. (2006), Bose and Daripab (2009), L.Bodoh-Creed (2012), Bose and Renou (2014), Wolitzky (2016), Song (2018), Lopomo et al. (2020), etc. In those papers, agents have ambiguous or misspecified beliefs over others' types or the underlying states. The ambiguity comes from either exogenous assumptions or endogenous ambiguous communication devices (Bose and Renou, 2014). Guo (2019) studies full surplus extraction by allowing for ambiguity in transfer schemes, and the characterization relies on correlated type distributions among agents. As a result, the framework in those papers does not naturally include the single-agent case, over which our framework still has some leverage.

Our approach parallels with Di Tillio et al. (2017) in the sense that the MD can endogenously introduce ambiguity in both allocation rules and transfer schemes. However, the focuses are distinct. Di Tillio et al. (2017) consider the revenue maximization problem in the context of selling an object to ambiguity averse buyers, while our chapter explores the implementation problem of a given choice correspondence. Apparently, the two works are complementary to each other.

The rest of the chapter is organized as follows. We formally set up the model in Section 2.2. We provide our main characterization result for the single-agent case in Section 2.3 and the multi-agent case in Section 2.4. In Section 2.5, we consider specialized settings of supermodular environments and one-dimensional type spaces. We then give an application of our result in Section 2.6. We discuss our results in Section 2.7. All omitted proofs are in Appendix B.

2.2 Model

We start with single-agent implementation. The case with multiple agents will be discussed in Section 2.4. We consider a scenario where an agent has private information and can report a message to the MD. Based on the reported message, the MD chooses an outcome and pays (charges) a transfer to (from) the agent. The MD wants to choose the best outcome and does not care about the transfer. Throughout the chapter, we use $\Delta(S)$ to denote the set of probability measures over a measurable space S. For any $\mu \in \Delta(S)$ and any measurable set $E \subseteq S$, the probability of E is denoted by $\mu[E]$. All σ-algebras are omitted.

The outcome space is a nonempty measurable space X. Generic elements of X are denoted by x, y, z, etc. The type space of the agent is a nonempty measurable space Θ, with generic elements $\theta, \theta', \hat{\theta}$, etc. We require the singleton set $\{\theta\}$ to be measurable for any $\theta \in \Theta$. The utility of the agent is type-dependent and quasi-linear, which is given by a measurable function $w : \Theta \times X \times \mathbb{R} \to \mathbb{R}$ such that $w(\theta, x, t) = u(\theta, x) - t$. $u(\theta, x)$ is the utility received by a type θ agent from outcome x, and t is the transfer paid by the agent. The agent is a MEU maximizer, i.e., for a set of distributions over outcomes and transfers $\Lambda \subseteq \Delta(X \times \mathbb{R})$, a type θ agent's payoff is given by

$$\inf_{\lambda \in \Lambda} \mathbb{E}_\lambda[u(\theta, x) - t].$$

A choice correspondence is a map $F : \Theta \rightrightarrows X$, where $F(\theta)$ is nonempty and measurable for each $\theta \in \Theta$. The MD aims to implement her target choice correspondence F using *mechanisms*. Specifically, a single-mechanism is a tuple (g, t) where $g : \Theta \to \Delta(X)$ is an allocation rule and $t : \Theta \to R$ is a transfer scheme. Under the single-mechanism (g, t), when the agent reports $\theta \in \Theta$, the outcome is chosen according to the distribution $g(\theta)$ and the transfer paid by the agent is $t(\theta)$. A multi-mechanism is a nonempty set of single-mechanisms $(g_i, t_i)_{i \in I}$, where the MD commits

to use some single-mechanism indexed by $i \in I$ but conceals which one to use. For any $j \in I$, (g_j, t_j) is said to be a single-mechanism of the multi-mechanism $(g_i, t_i)_{i \in I}$. We define the implementability of a choice correspondence as follows.

Definition 2.2.1 *A multi-mechanism $(g_i, t_i)_{i \in I}$ implements a choice correspondence F if*

1. *(Truth-telling) For any $\theta \in \Theta$ and $\beta \in \Delta(\Theta)$,*

$$\inf_{i \in I} \left\{ \int_X u(\theta, x) g_i(\theta)[dx] - t_i(\theta) \right\}$$

$$\geq \inf_{i \in I} \left\{ \int_\Theta \int_X u(\theta, x) g_i(\theta')[dx] \beta[d\theta'] - \int_\Theta t_i(\theta') \beta[d\theta'] \right\}.$$

2. *(Non-triviality) For any $\theta \in \Theta$,*

$$\inf_{i \in I} \left\{ \int_X u(\theta, x) g_i(\theta)[dx] - t_i(\theta) \right\} > -\infty.$$

3. *(Consistency) For any $i \in I$ and $\theta \in \Theta$, $g_i(\theta)[F(\theta)] = 1$.*

A choice correspondence F is implementable if there is a multi-mechanism implementing it.

Truth-telling says that it is optimal for the agent to report his true type. Since the agent is a MEU maximizer, a randomized report can possibly be strictly better than any deterministic report.[7] Allowing for randomized reports largely expands the set of deviating strategies of the agent. The truth-telling condition requires that all such deviations need to be excluded. Nevertheless, as we will show in Section 2.3, the set of implementable correspondences remains the same whether or not randomization is allowed.

Non-triviality imposes no restriction on the lower bound of the agent's utility from truth-telling, so long as there is one. This condition is a technical one. If the

[7]Note that maxmin expected utility is concave. Thus, an agent might have incentive to hedge.

payoff of the agent is allowed to be negative infinity, then any choice correspondence is implementable. To see this, consider a sequence of single-mechanisms such that the transfers charged by the MD go to positive infinity uniformly for all reports. As a result, the agent always receives a payoff of negative infinity and has no incentive to misreport.

Consistency says that the MD can always implement her desired outcomes no matter which single-mechanism she uses. As we have illustrated in the introduction, consistency indicates that the MD has no incentive to cherry-pick a particular single-mechanism of the multi-mechanism. This validates the MD's adoption of ambiguity. Moreover, the consistency condition also guarantees the credibility of the MD's commitment. As a result, the multi-mechanism is credible even if it is not implemented by a third party who has no conflicting interests. In Section 2.7.2, we provide more discussions on the commitment power of the MD when transfers matter.

Note that we restrict to direct mechanisms, where the agent's message space is equal to his type space. By the revelation principle, it is without loss of generality. The detailed proof is in the appendix. An implicit assumption here is that the message space for each single-mechanism of the multi-mechanism is the same. This assumption is necessary since otherwise the agent is able to know which single-mechanism is used once he knows the set of messages he can report.

2.3 Main Result

In this section, we characterize implementable choice correspondences. Our result relies on a simple assumption imposed on the given choice correspondence F.

Definition 2.3.1 *A choice correspondence F is bounded if for any $\theta, \theta' \in \Theta$, $\{u(\theta, x) : x \in F(\theta')\}$ is bounded.*

To state our condition, we define $\theta_{n+1} := \theta_1$ for any nonempty finite sequence of types $\{\theta_1, ..., \theta_n\} \subseteq \Theta$. F is said to satisfy *cyclical monotonicity* if for any nonempty finite sequence $\{\theta_1, ..., \theta_n\} \subseteq \Theta$,

$$\sum_{k=1}^{n} \sup_{x \in F(\theta_{k+1})} [u(\theta_{k+1}, x) - u(\theta_k, x)] \geq 0. \tag{2.1}$$

Theorem 2.3.2 *A bounded choice correspondence is implementable if and only if it satisfies cyclical monotonicity.*

Note that a choice correspondence F is bounded if it is a choice function, i.e., if $F(\theta)$ is a singleton set for each $\theta \in \Theta$. In this case, our theorem reduces to the characterization of implementable choice functions, which is given by Rockafellar (1970) and Rochet (1987).

We illustrate the necessity part of Theorem 2.3.2 for the case with two types. Suppose that the choice correspondence F can be implemented by a multi-mechanism $(g_i, t_i)_{i \in I}$. Consider a pair of types $\theta_1, \theta_2 \in \Theta$. For any $l, h \in \{1, 2\}$, let $V(\theta_l, \theta_h, i)$ denote the payoff of a type θ_l agent under the single-mechanism indexed by $i \in I$ when he reports his type to be θ_h. We have

$$V(\theta_l, \theta_h, i) = \int_X u(\theta_l, x) g_i(\theta_h)[dx] - t_i(\theta_h).$$

To simplify the analysis, assume that there exists a single-mechanism indexed by $i_{l \to h}$ achieving the infimum of $\inf_{i \in I} V(\theta_l, \theta_h, i)$. Therefore, $V(\theta_l, \theta_h, i_{l \to h})$ is the payoff of a type θ_l agent when he reports θ_h.

We now apply the implementability conditions. By truth-telling, the agent has no incentive to misreport:

$$V(\theta_1, \theta_1, i_{1 \to 1}) \geq V(\theta_1, \theta_2, i_{1 \to 2}), \tag{2.2}$$

$$V(\theta_2, \theta_2, i_{2\to 2}) \geq V(\theta_2, \theta_1, i_{2\to 1}). \tag{2.3}$$

The choice of $i_{l\to h}$ ensures that

$$V(\theta_1, \theta_1, i_{2\to 1}) \geq V(\theta_1, \theta_1, i_{1\to 1}), \tag{2.4}$$

$$V(\theta_2, \theta_2, i_{1\to 2}) \geq V(\theta_2, \theta_2, i_{2\to 2}). \tag{2.5}$$

Combing inequalities (2.2)(2.4) and (2.3)(2.5), we have

$$V(\theta_1, \theta_1, i_{2\to 1}) \geq V(\theta_1, \theta_2, i_{1\to 2}), \tag{2.6}$$

$$V(\theta_2, \theta_2, i_{1\to 2}) \geq V(\theta_2, \theta_1, i_{2\to 1}). \tag{2.7}$$

By summing up inequalities (2.6) and (2.7), we can eliminate the transfer terms. Reorganization gives

$$\int_X [u(\theta_1, x) - u(\theta_2, x)] g_{i_{2\to 1}}(\theta_1)[dx] + \int_X [u(\theta_2, x) - u(\theta_1, x)] g_{i_{1\to 2}}(\theta_2)[dx] \geq 0. \tag{2.8}$$

By consistency, we know for each type θ and each single-mechanism indexed by $i \in I$, $g_i(\theta)[F(\theta)] = 1$. Thus, condition (2.8) implies

$$\sup_{x\in F(\theta_1)} [u(\theta_1, x) - u(\theta_2, x)] + \sup_{x\in F(\theta_2)} [u(\theta_2, x) - u(\theta_1, x)] \geq 0,$$

which is exactly the cyclical monotonicity condition in (2.1) for the sequence $\{\theta_1, \theta_2\}$. For sequences containing more than two types, a similar argument can show that the cyclical monotonicity condition holds. Note that non-triviality is not applied for the above argument since we assume the infimum $\inf_{i\in I} V(\theta_l, \theta_h, i)$ to be achieved by some single-mechanism. Otherwise, by applying the non-triviality condition, we can use a limiting argument and derive the same conclusion.

We propose a constructive proof for the sufficiency part of Theorem 2.3.2. In order to highlight some important features of the multi-mechanism that we construct, we consider a simplified setting where for each pair of types θ and θ', the supremum $\sup_{x\in F(\theta')}\{u(\theta',x)-u(\theta,x)\}$ is achieved. The proof for the general case is contained in Appendix B.

Proposition 2.3.3 *Suppose that cyclical monotonicity holds for a bounded choice correspondence F, and for any pair of types θ and θ',*

$$\operatorname*{argmax}_{x\in F(\theta')}[u(\theta',x)-u(\theta,x)]\neq\emptyset.$$

Then for any selection

$$x_{\theta,\theta'}\in\operatorname*{argmax}_{x\in F(\theta')}[u(\theta',x)-u(\theta,x)],$$

there exists a multi-mechanism $(g_\theta,t_\theta)_{\theta\in\Theta}$ implementing F such that for any θ and $\theta'\in\Theta$, $g_\theta(\theta')=x_{\theta,\theta'}$ and

$$u(\theta,g_\theta(\theta))-t_\theta(\theta)=u(\theta,g_{\theta'}(\theta))-t_{\theta'}(\theta)\geq u(\theta,g_\theta(\theta'))-t_\theta(\theta'). \tag{2.9}$$

Proof of Proposition 2.3.3. We first argue that condition (2.9) is sufficient for F to be implemented by the multi-mechanism $(g_\theta,t_\theta)_{\theta\in\Theta}$. Consistency is trivially satisfied since $x_{\theta,\theta'}\in F(\theta')$. Non-triviality holds since the equality of condition (2.9) implies that the payoff by reporting θ for a type θ agent is

$$\inf_{\hat{\theta}\in\Theta}\{u(\theta,g_{\hat{\theta}}(\theta))-t_{\hat{\theta}}(\theta)\}=u(\theta,g_\theta(\theta))-t_\theta(\theta)>-\infty.$$

Finally, we check the truth-telling condition. Given any $\theta \in \Theta$ and $\beta \in \Delta(\Theta)$, the payoff of reporting according to β for a type θ agent satisfies

$$\begin{aligned}\inf_{\hat{\theta}\in\Theta}\left\{\int_{\bar{\theta}\in\Theta}[u(\theta, g_{\hat{\theta}}(\bar{\theta})) - t_{\hat{\theta}}(\bar{\theta})]\beta[d\bar{\theta}]\right\} &\leq \int_{\bar{\theta}\in\Theta}[u(\theta, g_{\theta}(\bar{\theta})) - t_{\theta}(\bar{\theta})]\beta[d\bar{\theta}] \\ &\leq u(\theta, g_{\theta}(\theta)) - t_{\theta}(\theta),\end{aligned}$$

where the first inequality holds by the definition of infimum and the second inequality comes from the inequality part of condition (2.9). As a result, the multi-mechanism satisfies truth-telling.

Now it suffices to construct the multi-mechanism satisfying condition (2.9). Let $g_{\theta}(\theta') = x_{\theta,\theta'}$ for any $\theta, \theta' \in \Theta$. What remains to be constructed is the transfer scheme. Define

$$N(\theta, \theta') := u(\theta', x_{\theta,\theta'}) - u(\theta, x_{\theta,\theta'}),$$

$$D(\theta, \theta') := -N(\theta, \theta').$$

By cyclical monotonicity, for any nonempty finite sequence $\{\theta_1, ..., \theta_n\} \subseteq \Theta$,

$$\sum_{k=1}^{n} N(\theta_k, \theta_{k+1}) = -\sum_{k=1}^{n} D(\theta_k, \theta_{k+1}) \geq 0.$$

Equivalently, we have

$$\sum_{k=1}^{n-1} D(\theta_k, \theta_{k+1}) \leq N(\theta_n, \theta_1). \tag{2.10}$$

For any $n \geq 2$, let $\mathcal{S}_n(\theta, \theta')$ denote the set of all finite sequences $\{\theta_1, ..., \theta_n\}$ with length n where $\theta_1 = \theta$ and $\theta_n = \theta'$. By inequality (2.10), we have

$$\sup_{n\geq 2}\left\{\sup_{\{\theta_1,...,\theta_n\}\in\mathcal{S}_n(\theta,\theta')}\left[\sum_{k=1}^{n-1} D(\theta_k, \theta_{k+1})\right]\right\} \leq N(\theta', \theta). \tag{2.11}$$

By this, we can define for any $\theta, \theta' \in \Theta$,

$$H(\theta, \theta') := \sup_{n \geq 2} \left\{ \sup_{\{\theta_1, \ldots, \theta_n\} \in \mathcal{S}_n(\theta, \theta')} \left[\sum_{k=1}^{n-1} D(\theta_k, \theta_{k+1}) \right] \right\}.$$

Inequality (2.11) ensures that H takes real values. It is easy to see that for any $\theta, \theta', \theta'' \in \Theta$,

$$D(\theta, \theta') + H(\theta', \theta'') \leq H(\theta, \theta''). \tag{2.12}$$

With these inequalities, we construct the transfer scheme. Fix some $\theta^* \in \Theta$ and define $t_\theta(\theta') = u(\theta', x_{\theta,\theta'}) - H(\theta', \theta^*)$. Now we verify condition (2.9). First, for any $\theta, \theta' \in \Theta$,

$$u(\theta, g_{\theta'}(\theta)) - t_{\theta'}(\theta) = u(\theta, x_{\theta',\theta}) - [u(\theta, x_{\theta',\theta}) - H(\theta, \theta^*)] = H(\theta, \theta^*),$$

which is independent of the index θ' of the single-mechanism. Thus, the equality part of condition (2.9) holds. For the inequality part, given any $\theta, \theta' \in \Theta$, we have

$$\begin{aligned} u(\theta, g_\theta(\theta')) - t_\theta(\theta') &= u(\theta, x_{\theta,\theta'}) - [u(\theta', x_{\theta,\theta'}) - H(\theta', \theta^*)] \\ &= D(\theta, \theta') + H(\theta', \theta^*) \leq H(\theta, \theta^*). \end{aligned}$$

The last inequality is given by condition (2.12). ■

Several features of the multi-mechanism constructed in Proposition 2.3.3 should be noted. First, when reporting truthfully, a type θ agent gets the same payoff under every single-mechanism of the multi-mechanism. The maxmin mechanism constructed by Di Tillio et al. (2017) also shares this feature. To see why this feature is desirable for the MD, consider a type θ agent and two single-mechanisms (g_i, t_i) and (g_j, t_j) where the single-mechanism (g_i, t_i) yields a strictly higher payoff for the agent when he is reporting θ than (g_j, t_j). In this case, the MD can increase the transfer $t_i(\theta)$ until the two single-mechanisms give the same payoff for the agent when he reports θ.

Such an increase of transfer not only maintains the incentive compatibility of the type θ agent, but also weakens the incentive of an agent of a different type to misreport θ.

Second, by condition (2.9), a type θ agent's incentive to tell the truth is ensured by the single-mechanism (g_θ, t_θ). Any misreport of a type θ agent yields him a weakly lower payoff under the single-mechanism (g_θ, t_θ) than his truth-telling payoff. By our construction, the single-mechanism (g_θ, t_θ) is an inferior single-mechanism for a type θ agent, which prevents him from hedging since the agent's payoff under this particular single-mechanism is linear in the distributions of his reports. Despite the debate in the literature about how decision makers randomize to eliminate ambiguity[8], our multi-mechanism avoids such an issue, and can be applied without worrying about whether the agent is able to randomize or not.

Third, each single-mechanism of the multi-mechanism is deterministic, i.e., each reported type is mapped to a certain outcome under each single-mechanism. This property also holds for the multi-mechanism we construct for the general case in the proof of Theorem 2.3.2. Thus, once the single-mechanism used is revealed, everything is deterministic and the MD does not need a randomization device.

Note that when randomized reports are not allowed, the necessity of cyclical monotonicity still holds since our arguments only make use of the incentive constraints for deterministic misreports. By this observation, we have the following corollary.

Corollary 2.3.4 *Suppose that the agent is not able to randomize his reports. A bounded choice correspondence is implementable if and only if it satisfies cyclical monotonicity.*

As a final remark, we discuss the following weakly undominated condition.

[8]See, for example, Saito (2015) and Ke and Zhang (2020).

(Weakly Undominated Condition) For each $\theta \in \Theta$, there does not exist any $\beta \in \Delta(\Theta)$ such that

$$\int_\Theta \int_X u(\theta, x) g_i(\theta')[dx]\beta[d\theta'] - \int_\Theta t_i(\theta')\beta[d\theta'] \geq \int_X u(\theta, x) g_i(\theta)[dx] - t_i(\theta)$$

for all $i \in I$ with at least one inequality being strict.

The weakly undominated condition says that for any $\theta \in \Theta$, truth-telling is not a weakly dominated strategy for a type θ agent. If this condition is violated, the agent might directly eliminate the truth-telling strategy, which impairs the implementation of the choice correspondence. We show in the appendix, following the proof of Theorem 2.3.2, that for any implementable choice correspondence, we can find a multi-mechanism that implements the choice correspondence and satisfies this condition.

2.4 Multiple Agents

In this section, we consider implementation with multiple agents. As we will demonstrate, the implementation problem with multiple agents can be decomposed into a set of implementation problems with one agent.

Let a nonempty finite set A denote the set of agents. The type space of agent $a \in A$ is a nonempty measurable space Θ^a, with each singleton set being measurable. Let $\Theta := \times_{a \in A} \Theta^a$ and $\Theta^{-a} := \times_{a' \in A: a' \neq a} \Theta^{a'}$. Θ is the space of type profiles of all agents, and Θ^{-a} is the space of type profiles of all agents but a. We will use θ^a, θ^{-a} and θ to denote generic elements in Θ^a, Θ^{-a} and Θ respectively. For each agent $a \in A$, $P^a \in \Delta(\Theta^a)$ is the prior distribution of agent a's types. We assume that agents' type distributions are independent. Thus, the prior is defined as $P := \times_{a \in A} P^a \in \Delta(\Theta)$, which is assumed to be common knowledge. We denote P^{-a} as the

marginal distribution of P over Θ^{-a}. The payoff of agent a is given by the function $u^a : \Theta \times X \to \mathbb{R}$. Note that our framework allows for interdependent valuations.

The MD designs mechanisms to implement a choice correspondence $F : \Theta \rightrightarrows X$. A single-mechanism is a tuple (g, t) where $g : \Theta \to \Delta(X)$ is the allocation rule, and $t : \Theta \to \mathbb{R}^A$ is the transfer scheme. Let t^a denote the ath coordinate of t, which is the transfer scheme for agent a. A multi-mechanism is a nonempty set of single-mechanisms $(g_i, t_i)_{i \in I}$. We have the following definition for implementability of a choice correspondence F.

Definition 2.4.1 *A multi-mechanism $(g_i, t_i)_{i \in I}$ implements a choice correspondence F if*

1. *(Truth-telling) For any $a \in A$, $\theta^a \in \Theta^a$ and $\beta \in \Delta(\Theta^a)$,*

$$\inf_{i \in I} \left\{ \int_{\Theta^{-a}} \left(\int_X u^a(\theta^a, \theta^{-a}, x) g_i(\theta^a, \theta^{-a})[dx] - t_i^a(\theta^a, \theta^{-a}) \right) P^{-a}[d\theta^{-a}] \right\}$$

$$\geq \inf_{i \in I} \left\{ \int_{\Theta^{-a}} \int_{\Theta^a} \left(\int_X u^a(\theta^a, \theta^{-a}, x) g_i(\hat{\theta}^a, \theta^{-a})[dx] - t_i^a(\hat{\theta}^a, \theta^{-a}) \right) \beta[d\hat{\theta}^a] P^{-a}[d\theta^{-a}] \right\}.$$

2. *(Non-triviality) For any $a \in A$ and $\theta^a \in \Theta^a$,*

$$\inf_{i \in I} \left\{ \int_{\Theta^{-a}} \left(\int_X u^a(\theta^a, \theta^{-a}, x) g_i(\theta^a, \theta^{-a})[dx] - t_i^a(\theta^a, \theta^{-a}) \right) P^{-a}[d\theta^{-a}] \right\} > -\infty.$$

3. *(Consistency) For any $i \in I$ and $\theta \in \Theta$, $g_i(\theta)[F(\theta)] = 1$.*

A choice correspondence F is implementable if there is a multi-mechanism implementing it.

The three conditions are analogous to the conditions stated in Definition 2.2.1 for the single-agent case. The only difference here is that truth-telling is not a dominant strategy for each agent. The solution concept we consider here is Bayesian-Nash equilibrium where agents assess the choice of each single-mechanism as being made before

the resolution of the uncertainty of other agents' types. We discuss two alternative implementation concepts in Section 2.7.1.

We argue that the implementation problem can be reduced to a set of interim implementation problems for each agent a. To see this, suppose that all agents except for a will report truthfully. The MD aims to incentivize agent a to tell the truth. In this case, we can transfer the problem to a single-agent implementation problem. Specifically, define $F^a : \Theta^a \rightrightarrows X^{\Theta^{-a}}$ such that

$$F^a(\theta^a) := \left\{\gamma^{-a} \in X^{\Theta^{-a}} : \gamma^{-a}(\theta^{-a}) \in F(\theta^a, \theta^{-a})\right\}.$$

To interpret, each function $\gamma^{-a} \in X^{\Theta^{-a}}$ in $F^a(\theta^a)$ is an outcome the MD wants to implement given that agent a's type is θ^a. This outcome further depends on other agents' types θ^{-a}, which are assumed to be truthfully reported and hence known by the MD. Under this setting, we can define agent a's utility function $\bar{u}^a$ over $X^{\Theta^{-a}}$ as

$$\bar{u}^a(\theta^a, \gamma^{-a}) := \int_{\Theta^{-a}} u^a(\theta^a, \theta^{-a}, \gamma^{-a}(\theta^{-a}))P^{-a}[d\theta^{-a}].$$

By this, we reduce the interim implementation problem to a standard implementation problem with one agent: the set of outcomes is $X^{\Theta^{-a}}$, the agent has utility function $\bar{u}^a : \Theta^a \times X^{\Theta^{-a}} \to \mathbb{R}$, and the choice correspondence is F^a. A similar boundedness assumption is needed.

Definition 2.4.2 *A choice correspondence F is bounded if for any agent $a \in A$ and any pair of types θ^a and $\hat{\theta}^a$, the set $\{u^a(\theta^a, \theta^{-a}, x) : \theta^{-a} \in \Theta^{-a}, x \in F(\hat{\theta}^a, \theta^{-a})\}$ is bounded.*

Clearly, F^a being implementable for each $a \in A$ under the utility function $\bar{u}^a$ is a necessary condition for F to be implementable. This suggests that F^a satisfies the cyclical monotonicity condition. Our next theorem asserts that for a bounded choice

correspondence F, F^a satisfying cyclical monotonicity condition for each $a \in A$ is also sufficient for F to be implementable.

Theorem 2.4.3 *A bounded choice correspondence F is implementable if and only if for each agent $a \in A$, F^a satisfies cyclical monotonicity condition under $\bar{u}^a$, i.e., for any nonempty finite sequence of types $\{\theta_1^a, ..., \theta_n^a\} \subseteq \Theta^a$,*

$$\sum_{k=1}^{n} \left\{ \int_{\Theta^{-a}} \left(\sup_{x \in F(\theta_{k+1}^a, \theta^{-a})} \left\{ u^a(\theta_{k+1}^a, \theta^{-a}, x) - u^a(\theta_k^a, \theta^{-a}, x) \right\} \right) P^{-a}[d\theta^{-a}] \right\} \geq 0.$$

We extend the implementability condition from the single-agent case to the multi-agent case in a similar way to Proposition 6.1 of Borgers (2015). The assumption of independent type distributions is essential for Theorem 2.4.3 to hold. For each single-mechanism, it guarantees that the probability distribution over outcomes and the expected transfer are the same for each type θ^a and for any other type $\hat{\theta}^a$ reporting θ^a.

2.5 Supermodularity

In this section, we study the implementability of choice correspondences in supermodular environments. For simplicity, we consider the single-agent case and maintain the notations used in Section 2.2. All results in this section can be easily extended to the multi-agent case based on our discussions in Section 2.4.

To start with, let $\succsim^t$ be a total order over the type space Θ. Let $\trianglerighteq$ be a preorder over the outcome space X.[9]

[9] A total order is a complete, transitive and anti-symmetric binary relation. A preorder is a reflexive and transitive binary relation. As is standard in the literature, we use $\succ^t$ and $\triangleright$ to denote the asymmetric parts of $\succsim^t$ and $\trianglerighteq$ respectively.

Definition 2.5.1 *u exhibits increasing differences (with respect to $\succsim^t$ and $\trianglerighteq$) if for $\theta \succ^t \theta'$ and $x \trianglerighteq y$,*

$$u(\theta, x) - u(\theta', x) \geq u(\theta, y) - u(\theta', y).$$

u exhibits strictly increasing differences (with respect to $\succsim^t$ and $\trianglerighteq$) if u exhibits increasing differences and for $\theta \succ^t \theta'$ and $x \triangleright y$, we have

$$u(\theta, x) - u(\theta', x) > u(\theta, y) - u(\theta', y).$$

Intuitively, $\theta \succ^t \theta'$ means that a type θ agent attaches larger marginal utility to higher outcomes than a type θ' agent for any two outcomes ranked by $\trianglerighteq$. We illustrate the above definition through the following Example 2.5.2.

Example 2.5.2 There are two divisible goods $i = 1, 2$. Let the type space of the agent be $\Theta = \{\theta, \theta', \theta'', \theta'''\} \subseteq \mathbb{R}^2_+$ where $\theta = (2, 1), \theta' = (1, 1)$, $\theta'' = (1, 2)$ and $\theta''' = (1, 3)$. For each $\theta^* \in \Theta$, θ^*_i is the type θ^* agent's utility from having one unit of good $i \in \{1, 2\}$. The outcome space is $X = \mathbb{R}^2_+$. The outcome $x = (x_1, x_2)$ means that the agent obtains x_1 units of good 1 and x_2 units of good 2. A type θ^* agent's utility from outcome x is $u(\theta^*, x) = \theta^*_1 x_1 + \theta^*_2 x_2$.

Define the total order $\succsim^t$ over Θ as $\theta \succ^t \theta' \succ^t \theta'' \succ^t \theta'''$. Define the preorder $\trianglerighteq$ over X such that $x \trianglerighteq x'$ if and only if $x_1 \geq x'_1$ and $x_2 \leq x'_2$. We show that u exhibits increasing differences. For any $\theta^*, \theta^{**} \in \Theta$, if $\theta^* \succ^t \theta^{**}$ and $x \trianglerighteq x'$, we have

$$\theta^*_1(x_1 - x'_1) + \theta^*_2(x_2 - x'_2) \geq \theta^{**}_1(x_1 - x'_1) + \theta^{**}_2(x_2 - x'_2)$$

since $\theta^* \succ^t \theta^{**}$ implies that $\theta^*_1 \geq \theta^{**}_1$ and $\theta^*_2 \leq \theta^{**}_2$. However, u does not exhibit strictly increasing differences. To see this, consider outcomes $x \triangleright x'$ where $x = (1, 0)$ and $x' = (1, 1)$. We have $u(\theta, x) - u(\theta, x') = u(\theta', x) - u(\theta', x')$.

Next, we propose a sufficient condition for cyclical monotonicity in the supermodular environment. To motivate the condition, consider a cycle $\theta \to \theta' \to \theta'' \to \theta$ where $\theta \succ^t \theta' \succ^t \theta''$. Cyclical monotonicity requires that

$$\sup_{x \in F(\theta')} [u(\theta', x) - u(\theta, x)] + \sup_{x \in F(\theta'')} [u(\theta'', x) - u(\theta', x)] + \sup_{x \in F(\theta)} [u(\theta, x) - u(\theta'', x)] \geq 0.$$

To ensure the above inequality, we want to find outcomes $x_{\theta,\theta'} \in F(\theta')$, $x_{\theta',\theta''} \in F(\theta'')$ and $x_{\theta'',\theta} \in F(\theta)$ such that $u(\theta', x_{\theta,\theta'}) - u(\theta, x_{\theta,\theta'}), u(\theta'', x_{\theta',\theta''}) - u(\theta', x_{\theta',\theta''})$ and $u(\theta, x_{\theta'',\theta}) - u(\theta'', x_{\theta'',\theta})$ are large. By the increasing differences property and the order of the types $\theta \succ^t \theta' \succ^t \theta''$, we want $x_{\theta,\theta'}$ and $x_{\theta',\theta''}$ to be $\trianglerighteq$-low and $x_{\theta'',\theta}$ to be $\trianglerighteq$-high. In particular, it can be verified that the above inequality holds when $x_{\theta'',\theta} \trianglerighteq x_{\theta,\theta'}$ and $x_{\theta'',\theta} \trianglerighteq x_{\theta',\theta''}$. Intuitively, for $\theta \succ^t \theta'$, we want the selected outcome for upward deviations to θ (e.g. $x_{\theta'',\theta}$ for $\theta'' \to \theta$) to be $\trianglerighteq$-higher than the selected outcome for downward deviations to θ' (e.g. $x_{\theta,\theta'}$ for $\theta \to \theta'$) to be $\trianglerighteq$-low. This gives the following definition.

Definition 2.5.3 *F satisfies binary monotonicity if there is a selection of $\bar{x}_\theta, \underline{x}_\theta \in F(\theta)$ for each type θ such that for any two types θ_H and θ_L, if $\theta_H \succ^t \theta_L$, then $\bar{x}_{\theta_H} \trianglerighteq \underline{x}_{\theta_L}$.*

We give an example of a choice correspondence that satisfies binary monotonicity as follows.

Example 2.5.4 Following Example 2.5.2, consider the choice correspondence $F(\theta) = \{x\}$, $F(\theta') = \{x', y'\}$, $F(\theta'') = \{x'', y''\}$ and $F(\theta''') = y$ where $x = (4, 2), x' = (4, 5), y' = (5, 4), x'' = (3, 5), y'' = (2, 3)$ and $y = (1, 4)$. We argue that F satisfies binary monotonicity. Since $F(\theta)$ and $F(\theta''')$ are both singletons, we let $\bar{x}_\theta = \underline{x}_\theta = x$ and $\bar{x}_{\theta'''} = \underline{x}_{\theta'''} = y$. Clearly, $x \trianglerighteq y$. It remains to define $\bar{x}_{\theta'}, \underline{x}_{\theta'}, \bar{x}_{\theta''}$ and $\underline{x}_{\theta''}$. Binary monotonicity is equivalent to that $x \trianglerighteq \underline{x}_{\theta'}$, $x \trianglerighteq \underline{x}_{\theta''}$, $\bar{x}_{\theta'} \trianglerighteq y$, $\bar{x}_{\theta''} \trianglerighteq y$ and $\bar{x}_{\theta'} \trianglerighteq \underline{x}_{\theta''}$. These conditions are satisfied if we let $\bar{x}_{\theta'} = y'$, $\bar{x}_{\theta''} = y''$, $\underline{x}_{\theta'} = x'$ and $\underline{x}_{\theta''} = x''$.

We note that the choice correspondence in Example 2.5.4 cannot be implemented by a single-mechanism since no single-mechanism can guarantee the truth-telling constraints among types θ, θ' and θ'''.[10] The next proposition asserts that binary monotonicity is sufficient for the implementability of a choice correspondence. Hence, the choice correspondence in Example 2.5.4 can be implemented by a multi-mechanism.

Proposition 2.5.5 *If u exhibits increasing differences and F satisfies binary monotonicity, then F is implementable.*

Now suppose for some $\theta \in \Theta$, $F(\theta)$ contains two outcomes x_θ^h and x_θ^l such that $x_\theta^h \trianglerighteq x \trianglerighteq x_\theta^l$ for each $x \in F(\theta)$ (e.g., in Example 2.5.4, $F(\theta') = \{x', y'\}$ satisfies $y' \trianglerighteq x'$ but $F(\theta'') = \{x'', y''\}$ does not satisfy this condition). In this case, the selection of $\bar{x}_\theta$ and $\underline{x}_\theta$ from $F(\theta)$ for the binary monotonicity condition is clear: $\bar{x}_\theta = x_\theta^h$ and $\underline{x}_\theta = x_\theta^l$. Hence, if $\trianglerighteq$ is a weak order and each $F(\theta)$ contains a $\trianglerighteq$-maximal outcome x_θ^h and a $\trianglerighteq$-minimal outcome x_θ^l, binary monotonicity is equivalent to the condition that $\theta \succ^t \theta'$ implies $x_\theta^h \trianglerighteq x_{\theta'}^l$. By the increasing differences property, when $\theta \succ^t \theta'$, we have

$$\sup_{x \in F(\theta)} [u(\theta, x) - u(\theta', x)] = u(\theta, x_\theta^h) - u(\theta', x_\theta^h), \tag{2.13}$$

and

$$\sup_{x \in F(\theta')} [u(\theta', x) - u(\theta, x)] = u(\theta', x_{\theta'}^l) - u(\theta, x_{\theta'}^l). \tag{2.14}$$

Clearly, $x_\theta^h \trianglerighteq x_{\theta'}^l$ ensures that the sum of (2.13) and (2.14) is nonnegative, which is condition (2.1) for the binary cycle $\{\theta, \theta'\}$. If condition (2.1) holds for all binary cycles, the choice correspondence is said to satisfy *weak monotonicity* (or *2-monotonicity*). In fact, when $\trianglerighteq$ is complete and u exhibits strictly increasing differences, weak monotonicity also implies binary monotonicity. Thus, we have the following proposition.

[10] To see this, consider any lottery $\beta \in \Delta(\{x', y'\})$, which is the lottery assigned to the agent when he reports θ'. First, note that $\theta_1 > \theta_1'$ and $\theta_2 = \theta_2'$. Implementability requires that $x_1 \geq \beta(x') \cdot x_1' + \beta(y') \cdot y_1'$. Next, note that $\theta_1' = \theta_1'''$ and $\theta_2' < \theta_2'''$. Implementability requires that $y_2 \geq \beta(x') \cdot x_2' + \beta(y') \cdot y_2'$. Clearly, the two inequalities cannot be satisfied simultaneously.

Proposition 2.5.6 *Suppose that $\trianglerighteq$ is complete, u exhibits strictly increasing differences, and for each $\theta \in \Theta$, $F(\theta)$ contains a $\trianglerighteq$-maximal outcome $\bar{x}_\theta$ and a $\trianglerighteq$-minimal outcome $\underline{x}_\theta$. The following conditions are equivalent.*

(i) *F is implementable.*

(ii) *F satisfies binary monotonicity.*

(iii) *F satisfies weak monotonicity.*

As a final remark, we discuss how the supermodular environment is related to one-dimensional type spaces studied in Chapter 5.6 of Borgers (2015). Suppose that all primitive conditions in Proposition 2.5.6 hold. If for any two outcomes $x, y \in X$, both $x \trianglerighteq y$ and $y \trianglerighteq x$ hold, then any choice correspondence is implementable. Hence, we consider the nontrivial case that there exists $x_0, y_0 \in X$ with $x_0 \triangleright y_0$. Following the notion of Borgers (2015), $\succ^t$ is the induced order of $\trianglerighteq$ on the type space Θ in the sense that $\theta \succ^t \theta'$ if and only if

$$u(\theta, x) - u(\theta, x') > u(\theta', x) - u(\theta', x'), \quad \text{for all } x, x' \in X \text{ with } x \triangleright x',$$

and

$$u(\theta, x) - u(\theta, x') = u(\theta', x) - u(\theta', x'), \quad \text{for all } x, x' \in X \text{ with } x \trianglerighteq x' \text{ and } x' \trianglerighteq x.$$

In this case, Θ is *one-dimensional with respect to* $\trianglerighteq$. Hence, Proposition 2.5.6 provides a characterization of implementable choice correspondences for the case with one-dimensional type spaces and shows that the implementability is equivalent to weak monotonicity.

2.6 Application: Allocating a Single Good

In this section, we extend our analysis of Example 2.1.1 and study a general implementation problem of allocating a single good. We consider a MD who has a possibly incomplete preference over all possible types of agents and wants to allocate a good to an agent of her desired types. One prominent example is that the government needs to delegate the construction of a bridge to some private construction firm through public procurement. We fully characterize implementable choice correspondences generated by the MD's preferences by applying Proposition 2.5.6.

The set of agents is denoted as A, which is nonempty and finite. The outcome is the designated winner of the good and thus the outcome space is $X = A$. Each agent a has a private type $\theta^a \in \Theta^a$, where Θ^a is nonempty and finite. We maintain the notations Θ and Θ^{-a} as in Section 2.4.

The MD's preference over different types of agents is summarized by a preorder $\succsim_G$ over $\cup_{a\in A}\Theta^a$. We use $\succ_G$ to denote the asymmetric part of the preorder. Note that by allowing for incompleteness of the preference, we allow for the possibility that the MD cannot rank two agent types. The MD would like to allocate the good to any agent whose type is not $\succ_G$-dominated by any other agent's type. Hence the choice correspondence to be implemented is $F : \Theta \rightrightarrows A$ where

$$F(\theta) = \Big\{a \in A : \theta^b \not\succ_G \theta^a, \forall b \in A\Big\}.$$

For each agent a, there is a function $v^a : \Theta^a \to \mathbb{R}$ representing the agent's payoff of winning before the transfer occurs. For instance, $v^a(\theta^a)$ can be the construction cost of a type θ^a firm in the public procurement example. Each agent's net payoff of losing is normalized to 0. We further require that for any two different types $\theta^a, \hat{\theta}^a$ of agent a, $v^a(\theta^a) \neq v^a(\hat{\theta}^a)$.

For any $a \in A$, agent a's type follows a distribution $P^a \in \Delta(\Theta^a)$. We assume different agents' type distributions are independent and commonly known. P and P^{-a} are defined correspondingly. For any type $\theta^a \in \Theta^a$ and any preference $\succsim_G$ of the MD (with the induced choice correspondence F), define $\mathcal{H}^a_{\succsim_G}(\theta^a)$ and $\mathcal{L}^a_{\succsim_G}(\theta^a)$ as follows:

$$\mathcal{H}^a_{\succsim_G}(\theta^a) = \sum_{\theta^{-a} \in \Theta^{-a}} P^{-a} \mathbb{1}_{\{\theta^a \in F(\theta^a, \theta^{-a})\}},$$
$$\mathcal{L}^a_{\succsim_G}(\theta^a) = \sum_{\theta^{-a} \in \Theta^{-a}} P^{-a} \mathbb{1}_{\{\{\theta^a\} = F(\theta^a, \theta^{-a})\}}.$$

$\mathcal{H}^a_{\succsim_G}(\theta^a)$ denotes the probability that agent a's type θ^a is not $\succ_G$-dominated by any other agent's type. $\mathcal{L}^a_{\succsim_G}(\theta^a)$ is the probability that agent a's type θ^a $\succ_G$-dominates any other agent's type. By definition, $\mathcal{H}^a_{\succsim_G}(\theta^a) \geq \mathcal{L}^a_{\succsim_G}(\theta^a)$. Intuitively, $\mathcal{H}^a_{\succsim_G}(\theta^a)$ and $\mathcal{L}^a_{\succsim_G}(\theta^a)$ are respectively the maximal and minimal interim winning probabilities of agent a given his type being θ^a, when each agent reports truthfully and the MD chooses one of her most desired agents. We say a preference $\succsim_G$ of the MD is *implementable* if its induced choice correspondence is implementable. Implementable preferences can be fully characterized by a quasi-monotonicity condition.

Definition 2.6.1 *$\succsim_G$ satisfies quasi-monotonicity under the prior P if for any agent $a \in A$ and any two types $\theta^a, \hat{\theta}^a \in \Theta^a$, $v^a(\theta^a) > v^a(\hat{\theta}^a)$ implies $\mathcal{H}^a_{\succsim_G}(\theta^a) \geq \mathcal{L}^a_{\succsim_G}(\hat{\theta}^a)$.*

Theorem 2.6.2 *A preference $\succsim_G$ is implementable if and only if it satisfies quasi-monotonicity under the prior P.*

Theorem 2.6.2 is implied by Proposition 2.5.6. By Theorem 2.4.3, the implementation problem can be reduced to the implementation problems for each agent. For agent a, the type space Θ^a is totally ordered according to v^a. The outcome space for agent a is the set of all possible interim probabilities of winning the good, which is ordered by the Euclidean order. Therefore, we can apply Proposition 2.5.6.

Quasi-monotonicity is weaker than the standard monotonicity condition for implementation with a single-mechanism, which requires the interim winning probability to be weakly increasing with respect to v^a for each agent a. As is clear from Example 2.1.1, there are implementable choice correspondences for which no single-mechanism can ensure the monotonicity of the interim winning probabilities for all agents. If the MD is only allowed to use one single-mechanism, she needs to ensure the incentives for truth-telling of all types of agents. In contrast, by using a multi-mechanism, each single-mechanism of the multi-mechanism is used to incentivize one particular type of agents. As a result, the constraints for each single-mechanism become much weaker.

An immediate observation from Theorem 2.6.2 is that if the MD becomes more indecisive among agents' types, her preference is more likely to be implementable. Intuitively, when the MD's preference is not defined over more pairs of types, the MD's choice correspondence maps each type profile to a larger set of outcomes. Thus, the cyclical monotonicity condition is easier to be satisfied since the supremum takes weakly higher values. Correspondingly, indecisiveness of the MD also means that the MD can create more ambiguity by hiding the tie-breaking rules. Consequently, $\mathcal{H}^a_{\gtrsim_G}(\theta^a)$ becomes larger and $\mathcal{L}^a_{\gtrsim_G}(\theta^a)$ becomes smaller for each type θ^a.

We hasten to point out that our characterization of implementable preferences is computationally easy to check. The highest and lowest interim winning probabilities of each type of agents is independent of the choice of the multi-mechanism and could be pinned down directly from the distributions of agents' types and the MD's preference.

As a final remark of the model, we discuss the case where the MD's preference is not implementable. If there is a third party who has no conflicting interests, the MD can delegate a multi-mechanism, which violates the consistency condition, to the third party. That is, some single-mechanisms of the multi-mechanism implement her sub-optimal outcomes. The third party then adopts an arbitrary single-mechanism

of the multi-mechanism. This ensures the credibility of the multi-mechanism. Note that the MD can always expand her choice correspondence by adding her sub-optimal outcomes until the choice correspondence becomes implementable. By this, the MD identifies her sub-optimal implementable choice correspondence.

2.7 Discussion

2.7.1 Multiple Agents: Alternative Implementation Concepts

In this section, we discuss two alternative implementation concepts for the multi-agent case. First, we consider dominant-strategy implementation.

Definition 2.7.1 *A multi-mechanism $(g_i, t_i)_{i\in I}$ dominant-strategy implements a choice correspondence F if*

1. (Ex-post Truth-telling) For any $a \in A$, $\theta \in \Theta$ and $\beta \in \Delta(\Theta^a)$,

$$\inf_{i\in I}\left(\int_X u^a(\theta,x)g_i(\theta)[dx] - t_i^a(\theta)\right)$$

$$\geq \inf_{i\in I}\int_{\Theta^a}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\hat{\theta}^a,\theta^{-a})[dx] - t_i^a(\hat{\theta}^a,\theta^{-a})\right)\beta[d\hat{\theta}^a].$$

2. (Ex-post Non-triviality) For any $a \in A$ and $\theta \in \Theta$,

$$\inf_{i\in I}\left(\int_X u^a(\theta,x)g_i(\theta)[dx] - t_i^a(\theta)\right) > -\infty.$$

3. (Consistency) For any $i \in I$ and $\theta \in \Theta$, $g_i(\theta)[F(\theta)] = 1$.

A choice correspondence F dominant-strategy implementable if there is a multi-mechanism dominant-strategy implementing it.

Compared to Definition 2.4.1, we now require that truth-telling is a dominance strategy for each agent. Clearly, the implementation problem reduces to a set of implementation problems for each agent given the type profile of other agents. A direct corollary of Theorem 2.3.2 characterizes the set of dominant-strategy implementable choice correspondences.

Corollary 2.7.2 *A bounded choice correspondence F is dominant-strategy implementable if and only if for each $a \in A$, each $\theta^{-a} \in \Theta^{-a}$, and each nonempty finite sequence of types $\{\theta_1^a, ..., \theta_n^a\} \subseteq \Theta^a$,*

$$\sum_{k=1}^{n} \left(\sup_{x \in F(\theta_{k+1}^a, \theta^{-a})} \left\{ u^a(\theta_{k+1}^a, \theta^{-a}, x) - u^a(\theta_k^a, \theta^{-a}, x) \right\} \right) \geq 0.$$

Based on Theorem 2.4.3 and Corollary 2.7.2, any dominant-strategy implementable choice correspondence is also implementable.

Next, we consider another notion of Bayesian implementation. In Definition 2.4.1, we implicitly assume that agents assess the choice of the single-mechanism as being made *before* the resolution of the uncertainty of other agents' types. An alternative assumption is that agents assess the choice of the single-mechanism as being made *after* the resolution of the uncertainty of other agents' types. To distinguish this implementation concept from the one given by Definition 2.4.1, we call it η-implementation.

Definition 2.7.3 *A multi-mechanism $(g_i, t_i)_{i \in I}$ η-implements a choice correspondence F if*

1. *(η-Truth-telling) For any $a \in A$, $\theta^a \in \Theta^a$ and $\beta \in \Delta(\Theta^a)$,*

$$\int_{\Theta^{-a}} \left\{ \inf_{i \in I} \left(\int_X u^a(\theta^a, \theta^{-a}, x) g_i(\theta^a, \theta^{-a})[dx] - t_i^a(\theta^a, \theta^{-a}) \right) \right\} P^{-a}[d\theta^{-a}]$$

$$\geq \int_{\Theta^{-a}} \left\{ \inf_{i \in I} \int_{\Theta^a} \left(\int_X u^a(\theta^a, \theta^{-a}, x) g_i(\hat{\theta}^a, \theta^{-a})[dx] - t_i^a(\hat{\theta}^a, \theta^{-a}) \right) \beta[d\hat{\theta}^a] \right\} P^{-a}[d\theta^{-a}].$$

2. (η-Non-triviality) For any $a \in A$ and $\theta^a \in \Theta^a$,

$$\int_{\Theta^{-a}} \left\{ \inf_{i \in I} \left(\int_X u^a(\theta^a, \theta^{-a}, x) g_i(\theta^a, \theta^{-a})[dx] - t_i^a(\theta^a, \theta^{-a}) \right) \right\} P^{-a}[d\theta^{-a}] > -\infty.$$

3. (Consistency) For any $i \in I$ and $\theta \in \Theta$, $g_i(\theta)[F(\theta)] = 1$.

A choice correspondence F is η-implementable if there is a multi-mechanism η-implementing it.

The only difference between η-implementation and implementation is the order of determination of single-mechanisms ($\inf_{i \in I}$) and the resolution of uncertainty about other agents' types ($\int_{\Theta^{-a}}$). It is easy to show that the cyclical monotonicity condition in Theorem 2.4.3 is also necessary for η-implementation. In fact, η-implementation is equivalent to implementation.

Proposition 2.7.4 *A bounded choice correspondence F is η-implementable if and only if it is implementable.*

2.7.2 Discussion of Assumptions

In this section, we discuss the assumption that the MD does not care about the transfers. First, we show that by imposing a stronger boundedness condition, there will not be any need for the MD to go beyond utilizing bounded transfers. Second, we address the potential concern regarding the commitment power of the MD when different single-mechanisms yield different expected transfers.

To begin with, we give a sufficient condition under which the MD can use bounded transfers to implement the choice correspondence. The boundedness of the choice correspondence cannot guarantee the uniform boundedness of the transfers.[11] For this

[11] Consider a simple example where the type space of the agent and the outcome space are both the set of natural numbers $\mathbb{N}$. The agent's utility is given by $u(n_\theta, n_x) = n_\theta n_x$ if his type is n_θ and the outcome is n_x. The choice correspondence to be implemented is $F(n) = \{n\}$. Clearly, this choice correspondence is implementable. The transfers used by the MD have to be unbounded. To

purpose, we assume that u is bounded. This is satisfied when (i) Θ and X are finite, or (ii) Θ and X are compact and u is a continuous function. By the construction of the transfer scheme in the proofs of Proposition 2.3.3 and Theorem 2.3.2, this assumption indeed ensures uniformly bounded transfers. For the remaining part of this section, we assume that u is bounded, and we only focus on multi-mechanisms with uniformly bounded transfers.

Next, we discuss the potential concern that when different single-mechanisms yield different expected total transfers, the MD might have incentive to choose the single-mechanism that has the highest expected total transfer. Hence, we want to impose an additional restriction that each single-mechanism yields the same expected total transfer. The following two propositions show that our characterization results hold when either (i) there is a single agent, and the agent is not allowed to randomize his reports, or (ii) there are at least two agents. Let P be the common prior over the type space of agents.

Proposition 2.7.5 *Assume that there is a single agent, and the agent is not allowed to randomize his reports. Assume further that P has full support on Θ. If a multi-mechanism $(g_i, t_i)_{i\in I}$ implements F, then there exists a multi-mechanism $(\bar{g}_j, \bar{t}_j)_{j\in J}$ implementing F such that*

$$\inf_{i\in I}\left(\int_{\theta\in\Theta} t_i(\theta)P[d\theta]\right) = \int_{\theta\in\Theta} \bar{t}_j(\theta)P[d\theta],\ \forall j \in J.$$

Proposition 2.7.6 *Assume that there are at least two agents. If a multi-mechanism $(g_i, t_i)_{i\in I}$ implements F, then there exists a multi-mechanism $(\bar{g}_j, \bar{t}_j)_{j\in J}$ implementing*

see this, let t_n be the payment of the agent when the reported type is n. The incentive constraint for a type-1 agent to not report n is: $1 \cdot 1 - t_1 \geq 1 \cdot n - t_n$. Thus, $t_n \geq n - 1 + t_1$, which indicates that the transfers are unbounded.

F such that

$$\inf_{i\in I}\left(\sum_{a\in A}\int_{\theta\in\Theta} t_i^a(\theta)P[d\theta]\right)=\sum_{a\in A}\int_{\theta\in\Theta}\bar{t}_j^a(\theta)P[d\theta],\ \forall j\in J.$$

The result in Proposition 2.7.6 still holds if agents are not allowed to randomize their reports.

In both the single-agent and the multi-agent cases, we want to guarantee that each single-mechanism in the multi-mechanism induces the same expected total transfer by properly defining the transfer schemes. When there is one agent, the agent's expected transfer must be a constant across different single-mechanisms. However, with more than one agent, an agent's expected transfer can vary across different single-mechanisms. Therefore, we have a stronger result for the multi-agent case. The following example shows that for the single-agent case, allowing randomized reports and requiring each single-mechanism to generate the same expected transfer strictly shrink the set of implementable choice correspondences.

Example 2.7.7 An agent has three possible types $\{\theta_1,\theta_2,\theta_3\}$. The MD wants to implement the choice correspondence F such that $F(\theta_1)=\{x_1\}$, $F(\theta_2)=\{x_2\}$ and $F(\theta_3)=\{x_3,x_3'\}$. Let u be the agent's utility function:

$$u(\theta_3,x_3)=u(\theta_1,x_3)=-10<u(\theta_1,x_3')=0,$$

$$u(\theta_3,x_3')=u(\theta_2,x_3')=-20<u(\theta_2,x_3)=0,$$

$$u(\theta_1,x_1)=u(\theta_1,x_2)=u(\theta_2,x_1)=u(\theta_2,x_2)=u(\theta_3,x_1)=u(\theta_3,x_2)=0.$$

It can be verified that the above choice correspondence F satisfies cyclical monotonicity since for each pair θ and θ', we have

$$\sup_{x\in F(\theta)}[u(\theta,x)-u(\theta',x)]=0.$$

Without loss of generality, we consider a multi-mechanism $\{g_i, t_i\}_{i\in I}$ that contains finite single-mechanisms and implements F. Due to the tightness of the cyclical monotonicity condition, it can be easily verified that

$$\min_{i\in I} V(\theta, \theta, i) = \min_{i\in I} V(\theta, \theta', i),$$

for each $\theta, \theta' \in \Theta$.

Now, consider a type θ_1 agent who randomizes his reports such that each type θ_i is reported with probability $\frac{1}{3}$. Let V^* be his payoff by making such a randomization. By truth-telling, we have

$$\min_{i\in I} V(\theta_1, \theta_1, i) \geq V^* \geq \frac{1}{3}\min_{i\in I} V(\theta_1, \theta_1, i) + \frac{1}{3}\min_{i\in I} V(\theta_1, \theta_2, i) + \frac{1}{3}\min_{i\in I} V(\theta_1, \theta_3, i).$$

It implies that

$$V^* = \frac{1}{3}\min_{i\in I} V(\theta_1, \theta_1, i) + \frac{1}{3}\min_{i\in I} V(\theta_1, \theta_2, i) + \frac{1}{3}\min_{i\in I} V(\theta_1, \theta_3, i).$$

Hence, there must exist a single-mechanism i that achieves $\min_{i'\in I} V(\theta_1, \theta_1, i')$, $\min_{i'\in I} V(\theta_1, \theta_2, i')$ and $\min_{i'\in I} V(\theta_1, \theta_3, i')$ simultaneously. A simple argument further establishes that for the single-mechanism i, $g_i(\theta_3) = x_3$. Without loss of generality, suppose that $t_i(\theta_1) = 0$. We have $t_i(\theta_2) = 0$ and $t_i(\theta_3) = -10$.

By a similar argument, we can show that there exists a single-mechanism j which achieves $\min_{i'\in I} V(\theta_2, \theta_1, i')$, $\min_{i'\in I} V(\theta_2, \theta_2, i')$ and $\min_{i'\in I} V(\theta_2, \theta_3, i')$ simultaneously, and satisfies $g_j(\theta_3) = x_3'$, $t_j(\theta_1) = 0, t_j(\theta_2) = 0$ and $t_j(\theta_3) = -20$. Hence, as long as the MD believes that the probability of the agent being type θ_3 is positive, the single-mechanism j yields a strictly higher expected transfer than the single-mechanism i.

2.7.3 Weak Monotonicity

In this section, we discuss the connection between weak monotonicity and cyclical monotonicity in the single-agent case. Formally, a choice correspondence satisfies weak monotonicity if for any two types θ and θ',

$$\sup_{x \in F(\theta)} [u(\theta, x) - u(\theta', x)] + \sup_{y \in F(\theta')} [u(\theta', y) - u(\theta, y)] \geq 0.$$

Weak monotonicity is appealing as it is imposed on every pair of types which is considerably easier to check in practice.

When the choice correspondence reduces to a choice function, a branch of literature has shown that weak monotonicity is equivalent to cyclical monotonicity when the type space is either one-dimensional or rich. We have discussed the case for one-dimensional type spaces in Section 2.5. In the following we discuss rich type spaces. Following Saks and Yu (2005), we assume that the outcome set X is finite, and the type space Θ is a subset of $\mathbb{R}^X$. For any type $\theta \in \mathbb{R}^X$, the utility function is defined as $u(\theta, x) := \theta(x)$. Implicitly, it is assumed that no pair of distinct types have the same utility function. By a theorem in Roberts (1979), Gui et al. (2004) show that when the type space is $\mathbb{R}^X$, weak monotonicity is equivalent to cyclical monotonicity. Bikhchandani et al. (2006) prove that the equivalence holds for some order-consistent domains, including $\mathbb{R}^X_+$. Saks and Yu (2005) further extend the result to any convex type space in $\mathbb{R}^X$.

However, for a general choice correspondence F, richness of the domain might not ensure the equivalence between weak monotonicity and cyclical monotonicity. Note that weak monotonicity only requires that for any two types θ and θ',

$$\sup_{x \in F(\theta)} [u(\theta, x) - u(\theta', x)] + \sup_{y \in F(\theta')} [u(\theta', y) - u(\theta, y)] \geq 0.$$

If $F(\theta) = X$, then for any $\theta' \in \Theta$, weak monotonicity trivially holds for θ and θ'. Thus, if except for a finite number of types in Θ, any other type θ satisfies $F(\theta) = X$, then we essentially lose the power of richness: only finitely many types' weak monotonicity conditions remain informative. As a result, in the setting of choice correspondence, a rich domain cannot guarantee the equivalence between weak monotonicity and cyclical monotonicity.

Chapter 3

Asymptotic Learning with Ambiguous Information

3.1 Introduction

[1]Consider an agent who relies on multiple information sources to learn a payoff-relevant state. A voter may depend on poll results and advertising to find out a politician's competence and agenda, and an investor uses reports of different analysts to forecast the future returns of a stock. A common assumption in the literature is that the decision maker has beliefs about the quality of their information sources, and that these beliefs are correctly specified. In such cases, asymptotic learning is successful. Although these assumptions are reasonable, in many settings, forming beliefs might not be straightforward. For example, consider a prospective customer consulting online reviews before making a purchase decision. She may not have particular beliefs about the quality of each reviewer, because they are being consulted for the first time. Such settings are widespread, yet little is known about learning in these environments. This chapter addresses this gap.

We analyze asymptotic learning when the decision maker lacks particular beliefs about her information sources. We study a decision maker who estimates a state by minimizing a loss function. She observes monotonic transformations of multiple unbiased signals. The state and the signals are jointly normally distributed, but the decision maker does not know the signals' precisions, that is, the inverse of their variances. The decision maker is not probabilistically sophisticated; instead, she is ambiguous regarding the precision of each information source, and perceives them to lie in a bounded interval. Each assignment of precisions to information sources pins down a belief of the decision maker, a joint distribution over signals, and the state. Thus, an interval of perceived precisions induces a set of beliefs. We assume the decision maker updates her beliefs prior-by-prior. That is, upon observing information, she updates each belief according to Bayes' rule, obtaining a posterior about

[1]This chapter is based on the paper "Asymptotic Learning with Ambiguous Information" joint with Pëllumb Reshidi and João Thereze. The paper has been presented in Tsinghua Conference on Behavioral, Experimental and Theoretical Economics 2021.

the state. In doing so, ambiguity about precisions induces ambiguity about the state, in that the agent considers multiple posteriors. Finally, she takes a robust approach and evaluates the expected loss according to the worst case across all posteriors.

This setup encompasses a broad range of environments. By modeling observables as monotonic transformations of signals, we cover cases in which the decision maker may observe unbiased signals directly, the actions of other agents, and many other possibilities. Furthermore, our assumptions on both distributions and ambiguity attitudes are for the sake of tractability. Our main insights go through in much more general settings.[2]

Our first result shows the induced ambiguity over the state does not vanish asymptotically. That is, the posterior beliefs of the decision maker do not converge to the same distribution as the number of information sources grows. We characterize this asymptotic set of posteriors. As in standard Bayesian learning, the variance of each posterior converges to zero. However, different beliefs lead to different weighting of signals, and consequently to different posterior means. For example, for any realization of signals, the agent's belief set contains a belief that assigns higher precisions to signals with high realizations and lower precisions otherwise. In this case, the posterior mean converges to a relatively high value. Similarly, there is a belief that leads to a relatively low posterior mean. Considering the set of all agent's beliefs generates an interval of posterior means. The set of asymptotic posteriors is the set of Dirac measures over that interval. Importantly, this set is independent of the objective of the decision maker.

Our second result characterizes the decision maker's asymptotic estimate of the state. Her decision problem can be interpreted as a zero-sum game against nature. First, the decision maker receives information and chooses the estimate that minimizes her expected loss. Afterwards, nature chooses the precision of each source with the

[2]See Section 3.6 for a detailed discussion of these extensions.

aim of maximizing the agent's loss. In doing so, nature affects the agent's posterior distribution. We show that, asymptotically, this is equivalent to nature choosing posterior means in the interval described in the previous paragraph. If the agent chooses a relatively low value within the interval, nature will maximize her loss by choosing the highest value possible, and vice versa. To minimize the maximal loss, the agent chooses an estimate that makes nature indifferent between choosing the highest or the lowest value in the interval of posterior means.

We show that, in our setting, asymptotic learning typically fails. That is, the agent's estimate is not consistent. Thus, we complement the vast literature on model misspecification, which obtains similar results by assuming the true parameter values are not in the support of the decision maker's prior distribution. By contrast, we maintain the assumption that the true precisions are in the set of beliefs the agent deems possible. In fact, we show the agent's estimate is typically inconsistent *even* in cases in which a misspecified Bayesian decision maker would learn the truth.

These results have several implications. First, our model provides a foundation for disagreement among agents even when information is abundant. In particular, we show that two agents with the same prior beliefs, who evaluate estimation errors differently, could disagree significantly asymptotically. Note that a Bayesian decision maker's posterior belief converges to a Dirac measure. Hence, regardless of the loss function, her estimate will be equal to this value. By contrast, our agent's asymptotic estimate depends on the particular form of the loss function. Thus, ambiguity about the precision of information sources might rationalize disagreement even between informed experts who aim to find out the truth: for example, scientists with access to the same large dataset.

Second, we argue a small amount of ambiguity can significantly amplify the effect of model misspecification. Concretely, assume the agent observes the unbiased signals directly. In that case, a Bayesian decision maker estimates the state correctly

even when she holds misspecified beliefs about the precision of information sources. By contrast, we show that for any amount of prior ambiguity, however small, misspecified agents exist who, when facing this ambiguity, experience arbitrarily large losses and estimation errors. Thus, the interaction between ambiguity and model misspecification might prove to be a fruitful direction for future research.

Third, we show the decision maker can be worse off even if she perceives all of her information sources as more informative. Consider two decision problems, a and b, in which the agent directly observes unbiased signals but has different intervals of perceived precision. We show that even if the lowest precision in a is higher than the highest precision in b, the decision maker may be better off under b. To carry out this comparison, we study how the initial ambiguity on precisions maps into induced ambiguity on the state. In particular, we show that the interval of posterior means is determined by the ratio between the highest and lowest possible perceived precisions. Because the length of this interval pins down the agent's loss, her welfare is monotonic in this ratio, regardless of the level of perceived precisions.

Last, we consider an application whereby the decision maker learns from others' actions, instead of observing signals directly. An ambiguity-averse econometrician observes choices by Bayesian decision makers who attempt to estimate a payoff-relevant state given their private information. She aims to estimate this state but does not know the precision of their private signals. For a concrete example, consider a healthcare official assessing the prevalence of a disease in a region. She relies on hospital reports to do so, but is not sure about the quality of their data collection protocols. We show the econometrician generically fails to aggregate information. We characterize how she may over- or underreact to the information contained in the observed actions, as a function of her prior beliefs and the true level of precisions.

Related Literature. Our chapter follows the literature on learning under ambiguity. Epstein and Schneider (2007) introduce a framework where an agent seeking to learn the state of the world, lacks confidence in their information about the environment. They consider the MaxMin Expected Utility model (MEU) following Gilboa and Schmeidler (1989) and a general updating rule for ambiguity that encompasses both prior-by-prior (full Bayesian) updating (Pires, 2002) and maximum likelihood updating (Gilboa and Schmeidler, 1993). Epstein and Schneider (2008b) study an application to a financial market where the representative agent observes one signal with ambiguous precision, and updates her beliefs prior by prior. They show how this ambiguity affects reactions to information and the asset price. Follow-up papers extend these results by incorporating ambiguity on the mean of the signals, and by considering equilibrium portfolio choices as well as general utility functions (Illeditsch, 2011; Gollier, 2011; Condie and Ganguli, 2017). In this chapter, we consider a similar setup as Epstein and Schneider (2008b) but focus on whether ambiguity vanishes, and whether the agent can estimate the state correctly, when the number of signals she observes goes to infinity.

Of relevance is also the literature on single-agent misspecified learning, which is another possible driving force for the failure of asymptotic learning. In this literature, a misspecified agent typically has a prior that assigns probability 0 to (a neighborhood of) the true model. Berk (1966) and Shalizi (2009) show that with exogenous information, under mild conditions, the agent's beliefs converge, although not to the true state. Other works focused on settings where the signals can be affected by the actions of the agent and are hence endogenous. Nyarko (1991) and Fudenberg et al. (2017) provide examples in which the convergence of beliefs fails. Similar to our setup, Heidhues et al. (2019) consider the convergence of beliefs and actions with a Gaussian prior and signals. Frick et al. (2020), Esponda et al. (2019), and Fudenberg et al. (2020) focus on the convergence results in general models with finite actions.

Our chapter differs from the existing work in three ways. First, the agent in the misspecified learning literature is a Bayesian learner, whereas in our setup, the decision maker holds multiple beliefs and adopts prior-by-prior updating. Second, the decision maker in our model is not misspecified in the sense that the true model is contained in her set of priors. Third, we show that in our setting, even when information is exogenous, as in Berk (1966) and Shalizi (2009), the belief set diverges almost surely.

Our chapter also relates to the robust statistics literature (Huber, 2004). Roughly speaking, robust statistics are statistics that produce good performance even with deviations from assumptions on the data generation process. Cerreia-Vioglio et al. (2013) highlight the close relation between decision making under ambiguity, akin to the approach in this chapter, and robust statistics, and characterize conditions under which the two approaches are equivalent. However, the problems studied in the robust statistics literature typically differ from the one studied in this chapter. For instance, Giacomini and Kitagawa (2021) and Giacomini et al. (2019) propose new tools for Bayesian inference in set-identified models to reconcile the asymptotic disagreement between Bayesian and frequentist inferences. By contrast, our focus is on whether information aggregation is successful as the number of sources grows without bound. Even in cases of point-identified models, ambiguity does not vanish in our setup because the precisions of different information sources are allowed to be different. Finally, this result is in contrast to Marinacci (2002), where ambiguity vanishes because all observations are drawn with from the same ambiguous distribution.

3.2 Setup

A decision maker aims to learn the state of the world, $\theta \in \Theta := \mathbb{R}$, and has access to N information sources. Denote the set of information sources as $I =: \{1, ..., N\}$. The prior distribution P_0 of the state θ is a normal distribution $\mathcal{N}\left(\mu, \frac{1}{\rho_\mu}\right)$, where

$\rho_\mu > 0$ is the inverse of the variance. We call ρ_μ the *precision* of the prior. The prior is common knowledge among the decision maker and all information sources. Each information source $i \in I$ features a signal $s_i = \theta + \varepsilon_i$, where the noise ε_i is normally distributed with mean 0 and precision $\rho_i > 0$, that is, $\varepsilon_i \sim \mathcal{N}\left(0, \frac{1}{\rho_i}\right)$. We assume that the state and all noises are independent from each other given precisions.

We further assume the actual precisions of information sources to be drawn i.i.d. from some distribution function G on $[\underline{\rho}, \overline{\rho}]$ with $\overline{\rho} > \underline{\rho} > 0$. The decision maker in our model is ambiguous about the precisions of her information sources. In particular, she knows that the precision of each information source lies in $[\underline{\rho}, \overline{\rho}]$, but she cannot form a probabilistic belief about it. The decision maker can form conjectures about the precision of any information source i. We denote the decision maker's conjectured precision as $\hat{\rho}_i \in [\underline{\rho}, \overline{\rho}]$. Finally, note the decision maker is not misspecified, because she does not deem the actual precisions as impossible ex ante. This observation follows from the assumption that the actual precision ρ_i of information source i lies in the perceived precision set $[\underline{\rho}, \overline{\rho}]$.

As mentioned in the introduction, the decision maker might be unable to observe the realized signals of the information sources. Instead, the observable for each information source i is $\boldsymbol{a}(s_i, \rho_i)$, which depends on the the realized signal and the precision of information source i. We assume that for each precision ρ_i, the observable $\boldsymbol{a}(s_i, \rho_i)$ is invertible as a function of s_i. Denote the realized observable as a_i and the inverted function for signals as $\boldsymbol{s}^{\boldsymbol{a}}(a_i, \rho_i)$. Given the observable a_i and the conjectured precision $\hat{\rho}_i$, the decision maker's conjectured signal is $\hat{s}_i = \boldsymbol{s}^{\boldsymbol{a}}(a_i, \hat{\rho}_i)$, which might be different from the actual realized signal s_i. Moreover, conditional on the realized state θ, the actual observables are i.i.d. according to the distribution function F on $\mathbb{R}$ where

$$F(a) = \int_{[\underline{\rho}, \overline{\rho}]} F_\rho\Big(\boldsymbol{s}^{\boldsymbol{a}}(a, \rho)\Big) dG(\rho),$$

with $F_\rho \sim \mathcal{N}\left(\theta, \frac{1}{\rho}\right)$ for each $\rho \in [\underline{\rho}, \overline{\rho}]$. Later in this chapter, we will discuss several different observables. For instance, the unbiased signal sources might be directly observable — $\boldsymbol{a}(s_i, \rho_i) = s_i$. We also study the case in which the decision maker can observe estimates of Bayesian agents based on their common prior and private signals — $\boldsymbol{a}(s_i, \rho_i) = \frac{\rho_i s_i + \rho_\mu \mu}{\rho_i + \rho_\mu}$.

Belief Updating. Denote the profile of precisions as $\rho^N := (\rho_1, ..., \rho_N)$, the profile of conjectured precisions as $\hat{\rho}^N := (\hat{\rho}_1, ..., \hat{\rho}_N)$, and the profile of observables as $a^N := (a_1, ..., a_N)$, and for each $n \geq 1$ the set of distributions over $\mathbb{R}^n$ as $\Delta(\mathbb{R}^n)$. Following Epstein and Schneider (2007) and Epstein and Schneider (2008b), we define $L^{\boldsymbol{a}}(\hat{\rho}^N, \theta) \in \Delta(\mathbb{R}^n)$ as the likelihood function for the profile of observables, which is the conditional distribution for observables given conjectured precisions $\hat{\rho}^N$ and the realized state θ. Then the set of likelihood functions of the decision maker can be represented by $\mathcal{L}_N^{\boldsymbol{a}}$, where

$$\mathcal{L}_N^{\boldsymbol{a}} = \{L^{\boldsymbol{a}}(\hat{\rho}^N, \theta) \in \Delta(\mathbb{R}^N) : \hat{\rho}^N \in [\underline{\rho}, \overline{\rho}]^N,\ \theta \in \mathbb{R}\}.$$

Note that to calculate the likelihood function of observables, one can first calculate the likelihood function of signals, which is just a multivariate normal distribution with independent marginals, and then make use of the one-to-one mapping between signals and observables given the profile of conjectured precisions.

We assume the decision maker adopts full Bayesian updating (Pires, 2002) to derive posteriors using the prior P_0 and the set of likelihood functions $\mathcal{L}_N^{\boldsymbol{a}}$. In other words, given the realized profile of observables a^N, and a vector of conjectured precisions $\hat{\rho}^N$, the posterior over the states $P_N^{\boldsymbol{a}}(a^N, \hat{\rho}^N) \in \Delta(\mathbb{R})$ is obtained by applying Bayes' rule.[3] Then, the posteriors of the decision maker can be represented by the

[3]Under this assumption, for each conjectured precision, the decision maker updates as if she were certain the conjecture is correct. Alternatively, we could allow the decision maker to update

following set:

$$\mathbb{P}^{\boldsymbol{a}}(a^N) = \Big\{ P_N^{\boldsymbol{a}}(a^N, \hat{\rho}^N) \in \Delta(\mathbb{R}) : \hat{\rho}^N \in [\underline{\rho}, \overline{\rho}]^N \Big\}.$$

After observing the profile of observables, the decision maker chooses an estimate g of the state θ to minimize some loss function $u(g-\theta)$. We assume $u : \mathbb{R} \to \mathbb{R}$ is strictly convex and minimized at 0. Given multiple beliefs, the decision maker is a maxmin expected utility (MEU) maximizer following Gilboa and Schmeidler (1989), and she evaluates her estimate based on the worst possible belief. This preference might be a result of the decision maker being ambiguity averse, or the decision maker's intention to derive a robust upper bound for the expected loss. That is, the decision maker's objective is to minimize the maximal expected loss across all distributions in the set of posteriors. She picks an estimate g to solve the following min-max problem:

$$\min_{g} \max_{p \in \mathbb{P}^{\boldsymbol{a}}(a^N)} \Big\{ \mathbb{E}_p \big[u(g-\theta) \big] \Big\}.$$

To maintain tractability, we made several assumptions. In Section 3.6, we discuss how our results depend on these assumptions. We find our main insights continue to hold for general signal distributions, as well as for ambiguity attitudes beyond MEU. By contrast, we highlight that our belief-updating rule plays a crucial role in our results. In the rest of the chapter, we focus on the limiting case in which the number of information sources N goes to infinity.

3.3 Asymptotic Beliefs

In this section, we characterize how the agent's posterior set behaves as the number of observables grows large. In particular, recall that for any N, the agent observes a^N. Given a profile of conjectured precisions $\hat{\rho}^N$, his posterior belief is $P_N^{\boldsymbol{a}}(a^N, \hat{\rho}^N)$.

her beliefs given a conjectured non-degenerate distribution about the precision of each signal. Under such conjectures our qualitative results still go through, however expressions become cumbersome.

We are interested in the asymptotic behavior of the agent's posterior set: $\mathbb{P}^{\boldsymbol{a}}(a^N)$. Thus, we define the limit set of posteriors:

$$\mathbb{P}^{\boldsymbol{a}}_{\infty}(a) = \{P : \exists \hat{\rho} \in [\underline{\rho}, \overline{\rho}]^{\infty} \text{ s.t. } P = \lim_{N\to\infty} P^{\boldsymbol{a}}(a^N, \hat{\rho}^N)\}$$

Note $\mathbb{P}^{\boldsymbol{a}}_{\infty}(a)$ is defined as the set of limits of posteriors that can be generated by some profile of precisions. That definition is silent about which posterior beliefs converge. In fact, many non-converging sequences of posterior beliefs exist, but, as Section 4 will make clearer, these sequences are immaterial for our discussion.

To characterize this set, we start by interpreting the optimization problem described in Section 3.2 as a zero-sum game between the decision maker and nature. Under this interpretation, after signals are realized, the decision maker chooses an estimate for the state to minimize her loss function. Subsequently, with knowledge of the estimate, nature is free to choose, for each signal, any precision within the uncertainty set of the decision maker. The decision maker's objective is then to guarantee the lowest loss conditional on the fact that nature acts after her and to her detriment.

3.3.1 Quadratic Loss with Observable Signals

To help build intuition, and as a rough sketch of the proof of our more general results, we describe and partially analyze the special case in which the loss function is quadratic, whereas the observables are simply equal to the realized signal values; that is, $u(g-\theta) = (g-\theta)^2$ and $\boldsymbol{a}(s_i, \rho_i) = s_i$. In other words, the decision maker has access to N unbiased and normally distributed signals $s_i = \theta + \varepsilon_i$, with $\varepsilon_i \sim \mathcal{N}\left(0, \frac{1}{\rho_i}\right)$. Recall the true ρ_i are unknown to the decision maker, who entertains an interval of perceived precision $[\underline{\rho}, \overline{\rho}]$. The decision maker's objective is to minimize the maximal mean-squared errors across all distributions in the set of posteriors. We denote by s^N the vector of the N observed signals. She picks an estimate g to solve the following

problem:

$$\min_{g} \max_{p \in \mathbb{P}^s(s^N)} \left\{ \mathbb{E}_p \left[(g - \theta)^2 \right] \right\}.$$

Due to the properties of the quadratic loss function, the above optimization problem can be simplified to one that only depends on the conditional mean and variance of the state, which can be calculated in closed form due to the assumption of joint normality. Denote them as $\mathbb{E}[\theta|s^N, \hat{\rho}^N] = \frac{\hat{\rho}^N \cdot s^N + \rho_\mu \mu}{\hat{\rho}^N \cdot \mathbb{1}^N + \rho_\mu}$ and $\mathbb{V}[\theta|s^N, \hat{\rho}^N] = \left(1 - \frac{\hat{\rho}^N \cdot \mathbb{1}^N}{\hat{\rho}^N \cdot \mathbb{1}^N + \rho_\mu}\right) \frac{1}{\rho_\mu}$ respectively. Then, the objective of the decision maker becomes

$$\min_{g} \max_{\hat{\rho} \in [\underline{\rho}, \overline{\rho}]^N} \left\{ \left(g - \mathbb{E}[\theta|s^N, \hat{\rho}^N] \right)^2 + \mathbb{V}[\theta|s^N, \hat{\rho}^N] \right\}.$$

By changing the precision of each signal, nature affects both the squared bias and the variance. It determines variance by choosing the sum of precisions across signals, and, importantly, it affects the posterior mean by assigning different precisions to different signal realizations. From the definition of the posterior variance, we see that as long as each signal is somewhat informative ($\underline{\rho} > 0$), as the number of available signals N increases, the posterior variance converges to 0. Hence, the more signals the decision maker receives, the more nature focuses on affecting the decision maker's loss function via the squared bias. In the extreme case, in which $N \to \infty$, for any choice of precisions that nature may consider, the posterior variance is equal to 0 and nature utilizes the square bias as its only lever. We next characterize nature's behavior when $N \to \infty$, where nature's choice of what precision to attribute to which signal exclusively affects the posterior mean.

Definition 3.3.1 *An assignment $\hat{\rho} : \mathbb{R}^\infty \to [\underline{\rho}, \overline{\rho}]^\infty$ is order-preserving if $s_i \leq s_j \implies \hat{\rho}_i \leq \hat{\rho}_j$ for all $i, j \in \mathbb{N}$; and it is order-reversing if $s_i \leq s_j \implies \hat{\rho}_i \geq \hat{\rho}_j$, for all $i, j \in \mathbb{N}$. An assignment is a threshold assignment if it is order-preserving or order-reversing and $\text{Im}(\hat{\rho}) \in \left\{ \underline{\rho}, \overline{\rho} \right\}^\infty$*

Lemma 3.3.2 *Let* $\hat{\rho}^*$ *solve* $\max_{\hat{\rho}\in[\underline{\rho},\overline{\rho}]^\infty} \left(g - \mathbb{E}[\theta|s,\hat{\rho}]\right)^2$ *for some* $g \in \mathbb{R}$. *Under observable signals,* $\hat{\rho}^*$ *is a threshold assignment.*

Nature finds it optimal to assign precisions to signals to maximize the squared bias. Intuitively, the way to do so is to either maximize or minimize the posterior mean: if the decision maker's estimate g is relatively low, nature finds it optimal to maximize the posterior mean, and vice versa. The intuition for Lemma 3.3.2 can be derived by analyzing the expression of the posterior mean when N goes to infinity. In that case, given an observed empirical distribution of signals, H, the expression for the posterior mean can be written as: $\mathbb{E}[\theta|s,\hat{\rho}] = \frac{\int s\hat{\rho}(s)dF(s)}{\int \hat{\rho}(s)dF(s)}$. Consider nature's choice to maximize this expression, while keeping the same sum of precisions $\int \hat{\rho}(s)dF(s) = c \in [\underline{\rho},\overline{\rho}]$. Because c pins down the denominator of the expression for the posterior mean, nature chooses an assignment to maximize $\int s\hat{\rho}(s)dF(s)$. To do so, nature assigns high-valued signals high precisions and low-valued signals low precisions, thereby moving the posterior mean towards higher signal realizations. Using the extreme precisions $\overline{\rho}$ and $\underline{\rho}$ is the best way to do that, therefore justifying the optimality of threshold strategies. Naturally, an analogous strategy is optimal to minimize the posterior mean.

To summarize this example, as the number of signals goes to infinity, nature focuses on affecting the agent's bias by strategically assigning precisions to signal realizations. Asymptotically, this is the only way nature can affect the agent's loss, as variance goes to zero regardless of the decision maker's conjecture. Finally, nature can implement this bias-maximizing behavior applying threshold strategies: monotonic precision assignments that use only extreme precisions.

3.3.2 Ambiguity Does Not Vanish

When signals are observable and the loss function is quadratic, we argued in the previous section that (i) nature can restrict attention to threshold strategies, and (ii)

as the number of signals goes to infinity, the set of posteriors converges to a set of degenerate distributions. We now show these two insights generalize.

First, we provide sufficient conditions on observables that guarantee nature can still restrict attention to threshold strategies. Recall that under general observables the agent cannot observe the realized signal, but rather has to backtrack those signals based on their conjectured precision. In particular, given conjectures precision $\hat{\rho}_i$ and observable realization a_i, the agent believes the signal realization is $\hat{s}_i = \boldsymbol{s}^{\boldsymbol{a}}(a_i, \hat{\rho}_i)$. In contrast to the observable signals case, when associating a particular precision to an observable realization, the agent changes his interpretation about the signal realization. The following assumption ensures the effect of this association does not break the monotonicity between observables and inverted signals that is required for simple threshold strategies to be optimal.

Assumption 3.3.3 *Define the weighted inverted signal function $g(\rho, x) \equiv \rho \boldsymbol{s}^{\boldsymbol{a}}(x, \rho)$. g is affine in ρ and strictly supermodular.*

This assumption allows for a broad range of observables relevant in several economics applications. The two examples described in the setup — directly observable unbiased signals and observable estimates from Bayesians with private information — satisfy this assumption. In the context of financial markets, the demand of CARA investors for an asset with value θ also satisfies Assumption 3.3.3. In particular, when the unbiased signals are the investor's private information, their demand for the risky asset is $\boldsymbol{a}(s_i, \rho_i) = \frac{1}{\alpha}(\rho_i s_i + \rho_\mu \mu)$, where α is their absolute risk aversion.

Lemma 3.3.4 *Let $\hat{\rho}^*$ solve $\max_{\hat{\rho} \in [\underline{\rho}, \overline{\rho}]^\infty} \left(g - \mathbb{E}[\theta | a, \hat{\rho}]\right)^2$ for some $g \in \mathbb{R}$. Under Assumption 3.3.3, $\hat{\rho}^*$ is a threshold assignment.*

We now address the asymptotic behavior of posteriors. As previously discussed, each strategy of nature corresponds to a plausible belief in the agent's belief set. In the previous section, as the number of observables went to infinity, we argued that nature

loses the ability to influence the posterior variance, as aggregate information becomes infinitely precise. However, by assigning precisions to signals, nature could still affect the agent's bias. With a general loss functions, higher moments of the posterior distribution are payoff-relevant for the agent. Nevertheless, the above rationale is preserved: all moments but the posterior mean become irrelevant asymptotically, and, in the limit, nature can only command the interval of posterior means. As a consequence, the set of posterior beliefs converges to an interval of degenerate distributions regardless of the loss function.

Theorem 3.3.5 *Let Assumption 3.3.3 hold. Define:*

$$\overline{m} = \frac{\underline{\rho}\int_{-\infty}^{\overline{m}} \boldsymbol{s}^{\boldsymbol{a}}(x,\underline{\rho})dF(x) + \overline{\rho}\int_{\overline{m}}^{\infty} \boldsymbol{s}^{\boldsymbol{a}}(x,\overline{\rho})dF(x)}{\underline{\rho}F(\overline{m}) + \overline{\rho}\,(1-F(\overline{m}))},$$
$$\underline{m} = \frac{\overline{\rho}\int_{-\infty}^{\underline{m}} \boldsymbol{s}^{\boldsymbol{a}}(x,\overline{\rho})dF(x) + \underline{\rho}\int_{\underline{m}}^{\infty} \boldsymbol{s}^{\boldsymbol{a}}(x,\underline{\rho})dF(x)}{\overline{\rho}F(\underline{m}) + \underline{\rho}\,(1-F(\underline{m}))}.$$

Then, for almost all sequences, a, of realized observables,

1. *For all sequences* $\hat{\rho} \in [\underline{\rho},\overline{\rho}]^{\infty}$

$$\underline{m} \leq \lim_{N\to\infty} \inf \mathbb{E}_{P_N(a^N,\hat{\rho}^N)}[\theta|a^N,\hat{\rho}^N] \leq \lim_{N\to\infty} \sup \mathbb{E}_{P_N(a^N,\hat{\rho}^N)}[\theta|a^N,\hat{\rho}^N] \leq \overline{m}.$$

2. *The limit set of posteriors is a set of degenerate distributions independent of s:*

$$\mathbb{P}_{\infty}(s) = \{\delta_b : \underline{m} \leq b \leq \overline{m}\}.$$

Theorem 3.3.5 formalizes the observation above. It starts by establishing that, for any precision assignment, posterior means are bounded by two real numbers: $\underline{m}$, $\overline{m}$. These numbers formalize the notion of maximal and minimal posterior means that nature can achieve asymptotically. The second part of the theorem shows that any converging posterior approximates a degenerate distribution, and that distribution

may have any mean between the boundaries $\underline{m}$ and $\overline{m}$. Finally, Theorem 3.3.5 characterizes the values of these boundaries. For example, $\overline{m}$ is generated by the following strategy of nature: give the highest precision to signals higher than $\overline{m}$ and the lowest precision to values below it. By giving more weight to high signals, nature moves the posterior mean up. The highest such posterior mean is the fixed point $\overline{m}$.

The fundamental consequence of Theorem 3.3.5 is that induced ambiguity on the state does not vanish asymptotically. Rather, the agent still entertains a wide range of values for the state θ even when he has access to an arbitrarily large number of informative observables. This finding is in stark contrast to quantifiable risk. In fact, a secondary consequence of the result above is that quantifiable risk completely disappears even in our setting: all the limit posteriors are degenerate around their means. In the next section, we show how the presence of ambiguity in the limit set of posteriors affects the optimal estimate of the agent.

3.4 Asymptotic Estimate

In this section, we characterize the asymptotic behavior of the decision maker. In particular, we are interested in analyzing how ambiguity with regards to the decision maker's information sources affects her ability to correctly estimate the state as the number of observables increases. Recall that, for each realization of observables a^N, her estimate $g^*(a^N)$ minimizes her loss function, considering the worst-case posterior in $\mathbb{P}^{\boldsymbol{a}}(a^N)$. Because observables and loss functions are arbitrary, obtaining an explicit solution to $g^*(a^N)$ for finite N is not an easy task, which makes a direct attempt at characterizing the solution as intractable. Rather, we leverage on Theorem 3.3.5 to solve this problem. The main result of this section characterizes the asymptotic

estimate by showing the following limit exchange holds.

$$\lim_{N\to\infty} \arg\min_{g} \max_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} \mathbb{E}_p\Big[u(g-\theta)\Big] = \arg\min_{g} \max_{p\in\mathbb{P}_\infty(a)} \mathbb{E}_p\Big[u(g-\theta)\Big].$$

Theorem 3.3.5 states that, for almost all realizations of observables, $\mathbb{P}_\infty(a) = \{\delta_m : m \in [\underline{m}, \overline{m}]\}$. The limit swap above suggests that, as N grows, the optimal estimate converges to the estimate of an agent who does not know the mean of θ but wants to guarantee the minimal loss in the interval $[\underline{m}, \overline{m}]$. This observation greatly simplifies the characterization: the asymptotic behavior of the estimate is pinned down by an extremely simple optimization problem. In this problem, the agent only cares about how biased her estimate is in the worst-case scenario. Recall that her loss is larger the further from the true state her estimate is. If her estimate is too far from $\underline{m}$, she has a large utility loss in the worst case, in which the state is actually $\underline{m}$. A symmetric argument holds for $\overline{m}$. Therefore, she guarantees minimal loss by being indifferent between these two extreme possible values of the state. This intuition is formalized in the next result.

Theorem 3.4.1 *$g^*(s^N) \xrightarrow{a.s.} g^*$, where g^* is the unique solution to $u(g^* - \underline{m}) = u(g^* - \overline{m})$.*

Although intuitive, this result depends on the non-trivial exchange of limits mentioned above. A priori, it is not clear that the limit swap holds. First, limits of optimizers of a sequence of optimization problems are not guaranteed to coincide with the optimizers of the limit problem. Second, not all distributions in the set $\mathbb{P}^{\boldsymbol{a}}(a^N)$ converge. Indeed, sequences of precisions always exist such that posterior beliefs diverge. Still, our result confirms the limit exchange is valid and the heuristic argument we gave above goes through formally. We make this argument in two steps, addressing each of the concerns highlighted above.

The first step is to show the decision maker's optimization can be approximated by an optimization that considers only the mean of posterior distributions, as N grows large. For any finite N, the decision maker's loss is clearly affected by higher moments of the posteriors, but because quantifiable risk vanishes as the number of observable information sources grows, the mean progressively becomes the only relevant moment. The second step relies on an extension of the Glivenko-Cantelli theorem. It provides the important result that $g^*(a^N)$ are bounded. Recall, from part 1 of Theorem 3.3.5 that non-converging posteriors are bounded. Thus, intuitively, N by N, the payoff obtained by a non-converging sequence can be bounded by the payoff of two converging sequences, so that restricting attention to the converging ones turns out to be without loss of generality. As a consequence, non-converging beliefs are innocuous: we can characterize the asymptotic behavior of the agent's estimate without addressing them. We prove these two steps are sufficient to guarantee the convergence of g^*.

Theorem 3.4.1 shows the asymptotic estimate is typically incorrect. To illustrate, recall that $[\underline{m}, \overline{m}]$ in Theorem 3.3.5 are independent of the particular choice of the loss function. Rather, they are determined by the initial ambiguity and the observable function $\boldsymbol{a}$. By contrast, the asymptotic estimate is a consequence of the behavior of the loss function, u, on the interval of posterior means $[\underline{m}, \overline{m}]$. This finding suggests the decision maker estimates the state correctly asymptotically only in the knife-edge case in which her loss function coincides with the observable function in a particular way. Moreover, in that case, perturbing either of these functions would again lead to an incorrect limit estimate. This result is particularly striking when compared to the behavior of a Bayesian agent who knows the precision of each source.[4] Because observables map one to one to signals conditional on precisions, a Bayesian's asymp-

[4]We could also consider the comparison with a Bayesian agent who does not know the precision of each information source, but rather entertains a distribution over those precisions. The comparison remains the same as long as their statistical model is identified.

totic estimate would be equal to the state, regardless of loss function. The next two examples illustrate this result.

Example 3.4.2 (Back to Quadratic Loss with Observable Signals) We revisit the example from Section 3.3.1 to show a case in which the agent correctly estimates the state asymptotically. By Theorem 3.4.1, we have that, under quadratic losses, $g^* = \frac{\overline{m}+\underline{m}}{2}$. Because signals are observable, Theorem 3.3.5 states that $\underline{m}$ and $\overline{m}$ are defined by:

$$\overline{m} = \frac{\underline{\rho}\int_{-\infty}^{\overline{m}} x dF(x) + \overline{\rho}\int_{\overline{m}}^{\infty} x dF(x)}{\underline{\rho}F(\overline{m}) + \overline{\rho}\,(1 - F(\overline{m}))}, \quad \underline{m} = \frac{\overline{\rho}\int_{-\infty}^{\underline{m}} x dF(x) + \underline{\rho}\int_{\underline{m}}^{\infty} x dF(x)}{\overline{\rho}F(\underline{m}) + \underline{\rho}\,(1 - F(\underline{m}))}.$$

Normality implies the real distribution of signals F is symmetric around the true state θ. Thus, in this case, the decision maker estimates the state correctly. This result is a consequence of the symmetric loss function, as well as the symmetry of the normal distribution, and the assumption of observable signals.

Corollary 3.4.3 $g^*(s^N) \xrightarrow{a.s.} \theta$.

The next example breaks this coincidence to show this result is knife-edged.

Example 3.4.4 (Asymmetric Loss with Observable Signals) In this example, we maintain the assumption of observable signals, but let the loss function be given by:

$$u(g-\theta) = \begin{cases} (g-\theta)^2 & \text{if } \ g \geq \theta \\ \lambda(g-\theta)^2 & \text{if } \ g < \theta \end{cases}.$$

with $\lambda > 0$. That is, the decision maker's loss is different for over- and under-estimating the state θ. If $\lambda > 1$, for example, the agent is less concerned with losses when she overestimates the true state, compared to when she underestimates it. Her concern could be lower for many reasons. For instance, a health official who wants to learn about the prevalence of a disease in a population would likely be more affected

if they believes the transmission rate is lower than it really is than if they believe it is higher. Conversely, a product developer may face a much higher personal loss if they believe demand is higher than it actually is and end up developing a costly product that fails to be marketed.

Following Theorem 3.4.1, we have that the optimal estimate satisfies $g^* = \frac{\underline{m}+\sqrt{\lambda}\overline{m}}{1+\sqrt{\lambda}}$. However, because the results of Theorem 3.3.5 do not depend on the loss function, we still have that $\frac{\underline{m}+\overline{m}}{2} = \theta$. Thus, the agent estimates incorrectly for any $\lambda \neq 1$. In particular, if $\lambda < 1$, her optimal estimate is below the real value of the state: $g^* < \theta$. This example shows how an environment in which an agent estimates correctly can be easily perturbed so that the agent no longer estimates the state correctly. Note that, we maintain the assumption that signals are observable, but depart from the assumption of symmetric losses.

In the next section, we explore applications of our model. In one of these applications, we depart from Example 3.4.2 by changing the assumption of signal observability instead of the symmetry assumption. We characterize the optimal estimate for that case and show how, again, the decision maker generically fails to correctly estimate the state.

A Foundation for Disagreement. Example 3.4.4 also highlights that the loss function directly affects the agent's asymptotic estimate. Note a Bayesian decision maker's posterior belief converges to a Dirac measure on the real state. Thus, with multiple Bayesian agents, as the available information grows, regardless of their loss functions, Bayesian agents will agree on the optimal estimation of the state.[5] By contrast, our agent's asymptotic estimate continues to depend on the particular form of the loss function. Thus, ambiguity about the precision of information sources might rationalize disagreement even between informed experts who aim to find out the truth,

[5]With observable signals, this result holds true even for misspecified Bayesian agents who wrongly perceive the precision of the signals. As the amount of information grows without bounds, their estimates converge to the same value.

for example, scientists with access to the same large dataset. This discussion offers a foundation for disagreement among agents even when information is abundant.

3.5 Applications

In this section, we explore different applications of our main results by changing the set of observables available to the decision maker or her loss function.

3.5.1 Comparative Statics of Ambiguity

First, we revisit Example 3.4.2 on quadratic loss with observable signals. We claim that, contrary to intuition, making all signals more precise is not necessarily beneficial to the decision maker.

Recall that by Theorem 3.3.5, the limit set of posteriors is a set of degenerate distributions δ_b with $\underline{m} \leq b \leq \overline{m}$, where

$$\overline{m} = \frac{\underline{\rho} \int_{-\infty}^{\overline{m}} x dF(x) + \overline{\rho} \int_{\overline{m}}^{\infty} x dF(x)}{\underline{\rho} F(\overline{m}) + \overline{\rho}\,(1 - F(\overline{m}))}, \quad \underline{m} = \frac{\overline{\rho} \int_{-\infty}^{\underline{m}} x dF(x) + \underline{\rho} \int_{\underline{m}}^{\infty} x dF(x)}{\overline{\rho} F(\underline{m}) + \underline{\rho}\,(1 - F(\underline{m}))}.$$

To see how the set of precisions deemed possible by the decision maker affects the limit set of posterior beliefs, first note $\overline{m}$ and $\underline{m}$ only depend on the fraction of the highest and the lowest possible precisions, instead of their absolute values, because we can rewrite $\overline{m}$ and $\underline{m}$ as

$$\overline{m} = \frac{\int_{-\infty}^{\overline{m}} x dF(x) + \eta \int_{\overline{m}}^{\infty} x dF(x)}{F(\overline{m}) + \eta\,(1 - F(\overline{m}))}, \quad \underline{m} = \frac{\eta \int_{-\infty}^{\underline{m}} x dF(x) + \int_{\underline{m}}^{\infty} x dF(x)}{\eta F(\underline{m}) + (1 - F(\underline{m}))},$$

where $\eta = \frac{\overline{\rho}}{\underline{\rho}}$. The following proposition shows both $\overline{m}$ and $\underline{m}$ change with η monotonically.

Proposition 3.5.1 *Let $\eta = \frac{\overline{\rho}}{\underline{\rho}} \in (1, +\infty)$. Under observable signals, $\overline{m}$ is monotonically increasing in η and $\underline{m}$ is monotonically decreasing in η. Moreover, when $\eta \to +\infty$, we have $\overline{m} \to \infty$ and $\underline{m} \to -\infty$; when $\eta \to 1$, we have $\overline{m} - \underline{m} \to 0$.*

In Proposition 3.5.1, $\eta = \frac{\overline{\rho}}{\underline{\rho}}$ can be interpreted as the degree of ambiguity in the set of possible precisions $[\underline{\rho}, \overline{\rho}]$. When more ambiguity exists regarding precisions of signals ex ante, the limit set of posteriors also expands, and hence, ambiguity regarding states is greater ex post.

Now, we explore the welfare implication of such comparative statics. By Corollary 3.4.3 of Theorem 3.4.1, the decision maker always estimates correctly at the limit when signals are observable. Thus, the optimal utility depends solely on the size of the limit set of posterior means, that is, $\overline{m} - \underline{m}$. Corollary 3.5.2 directly follows from Proposition 3.5.1.

Corollary 3.5.2 *Under observable signals, as η increases, the decision maker is strictly worse off asymptotically.*

Corollary 3.5.2 has two possibly counter-intuitive implications. First, it implies that if we fix $\underline{\rho}$ and increase $\overline{\rho}$, the decision maker is strictly worse off. That is, she prefers all of her signals to be imprecise rather than being ambiguous that some signals might be more precise. Second, consider two decision problems with the set of possible precisions given by $[\underline{\rho}_1, \overline{\rho}_1]$ and $[\underline{\rho}_2, \overline{\rho}_2]$, respectively. If $\underline{\rho}_2 > \overline{\rho}_1$ and $\eta_1 = \frac{\overline{\rho}_1}{\underline{\rho}_1} < \eta_2 = \frac{\overline{\rho}_2}{\underline{\rho}_2}$, the decision maker believes that any signal in the second decision problem is more precise than any signal in the first one, but she is strictly worse off in the second decision problem.

This result shows that making all signals more precise is not necessarily beneficial to the decision maker.

3.5.2 Misspecification and Ambiguity

In this section, we show that a small amount of ambiguity might substantially amplify the effect of model misspecification. We say that an agent is misspecified if the real precisions of information sources are not in her consideration set, that is, $\operatorname{supp} G \cap [\underline{\rho}, \overline{\rho}] = \emptyset$, where G is the true distribution of precisions. To focus on a simple environment, we proceed under the assumption that the agent observes unbiased signals directly, but has an asymmetric loss function u as in Example 3.4.4.

Note the impact of prior ambiguity can still be captured by the ratio η defined in the last section. For the special case of no ambiguity, $\eta = 1$. We say an agent is a misspecified Bayesian, if she is misspecified and faces no ambiguity, $\eta = 1$.

Proposition 3.5.3 *Assume signals are observable and u is as in Example 3.4.4. For any $\eta > 1$, and any constant $C > 0$, true distributions of precisions G exist such that $|g^* - \theta| > C$.*

To see the connection of this proposition with misspecified Bayesians, note the misspecified Bayesian agent believes the precision of information source i to be $\hat{\rho}_i \notin \operatorname{supp} G$. However, because signals are unbiased, this misspecification has no effect on asymptotic learning, regardless of the distribution G. In particular, for any loss function, the agent estimates the state correctly and, therefore, experiences zero loss. Now, consider this same agent when facing a small amount of ambiguity, η, very close to 1. By Proposition 3.5.3, true distributions exist over precisions, G, such that this agent makes arbitrarily large estimation errors. Thus, any amount of ambiguity, no matter how small, might be sufficient to cause large deviations from the correct estimate. These large deviations occur in settings in which the true precision of the signals is rather low. This observation is inconsequential for the misspecified Bayesian agent, who still estimates the state correctly. However, in the presence of ambiguity, the less informative the signals are in reality, the bigger this gap will be.

3.5.3 Aggregating Estimates

Next, we study next the problem of an ambiguity-averse econometrician who aims to estimate the state by aggregating estimates by numerous Bayesian agents. Agents share the same prior but have access to different information sources. Although the econometrician knows the prior distribution of the state, she does not know the precision of the individual sources. This environment is reasonable in many applications. For instance, consider a healthcare official assessing the prevalence of a disease in a region. She relies on hospital reports to do so, but is not sure about the quality of their data-collection protocols.

Formally, we assume all agents and the decision maker share the same prior beliefs about the state θ. As in the previous sections, according to the prior, $\theta \sim \mathcal{N}\left(\mu, \frac{1}{\rho_\mu}\right)$. Conditional on the realization of θ, agent i receives a private signal $s_i = \theta + \varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}\left(0, \frac{1}{\rho_i}\right)$. That is, each agent receives an unbiased signal about the state. We consider the case in which each agent attempts to estimate the realized value of θ to minimize the mean-squared error. Given the prior and the private signal, the optimal Bayesian estimate for agent i would then be $\boldsymbol{a}(s_i, \rho_i) = \mathbb{E}[\theta|\rho_i, s_i] = \frac{\rho_\mu \mu + \rho_i s_i}{\rho_\mu + \rho_i}$. These actions are the ones the econometrician has access to. The setup studied in this section is graphically depicted in Figure 3.1.

Figure 3.1: Learning From Actions Setup

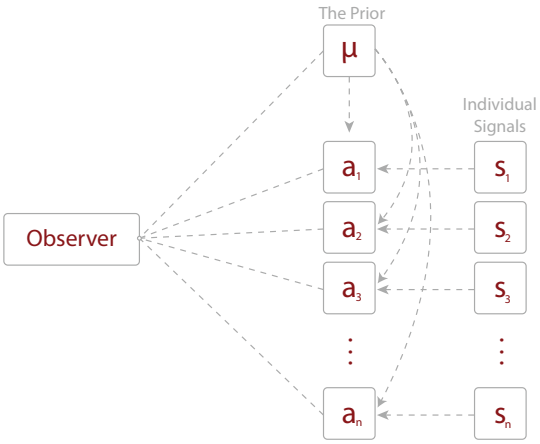

Although each agent knows the precision of their private signal, the econometrician does not. We once more assume that for each signal, the econometrician considers a set of possible precisions $\left[\underline{\rho}, \overline{\rho}\right]$. Because each action is a convex combination of the private signal s_i and the mean of the prior μ, an econometrician who intends to estimate the value of θ, will first have to transform the actions back to signals. For a conjectured precision $\hat{\rho}_i$, the recovered signal will be $\boldsymbol{s^a}(a_i, \hat{\rho}_i) = a_i + \frac{\rho_\mu}{\hat{\rho}_i}(a_i - \mu)$. We again assume the loss function of the econometrician is quadratic. We start by utilizing Theorem 3.3.5 to characterize the limit set of posteriors of the econometrician in this example.

Proposition 3.5.4 *Let* $\boldsymbol{a}(s_i, \rho_i) = \frac{\rho_\mu \mu + \rho_i s_i}{\rho_\mu + \rho_i}$. *The boundaries of the limit set of posteriors for the econometrician are:*

$$\overline{m}_a = \frac{\underline{\rho} \int_{-\infty}^{\overline{m}_a} x dF(x) + \bar{\rho} \int_{\overline{m}_a}^{\infty} x dF(x) + c}{\underline{\rho} F(\overline{m}_a) + \bar{\rho}\,(1 - F(\overline{m}_a))} \qquad \underline{m}_a = \frac{\overline{\rho} \int_{-\infty}^{\underline{m}_a} x dF(x) + \underline{\rho} \int_{\underline{m}_a}^{\infty} x dF(x) + c}{\overline{\rho} F(\underline{m}_a) + \underline{\rho}\,(1 - F(\underline{m}_a))},$$

where $c = (\theta - \mu) \int \frac{\rho \rho_\mu}{\rho_\mu + \rho} dG(\rho)$.

The boundaries of the limit set of posteriors are defined by a fixed point similar to the one from the case with observable signals. However, here, the econometrician has to backtrack realized signals from observed estimates, which leads to an adjustment term c. The next result is a corollary of Theorem 3.4.1.

Corollary 3.5.5 *For almost all sequences* a *and values of the state* θ, $\lim_{N \to \infty} |g^*(a^N) - \theta| > 0$.

That is, the econometrician's estimation converges away from the truth almost surely, because inverting from observables to signals depends on the conjectured precisions and of the prior mean. The lack of knowledge about the former makes distinguishing signal realizations from the prior mean impossible, thus generating a bias in the recovered signals. The following assumption allows us to clearly characterize the optimal estimate and to analyze comparative statics.

Assumption 3.5.6 *For some* $\rho^* \in [\underline{\rho}, \overline{\rho}]$, $G = \delta_{\rho^*}$.

Although the econometrician might consider different precisions for each signal, under Assumption 3.5.6, in reality, all signals share the same precision. This assumption allows us to characterize how the econometrician estimate differs from the true parameter value. We say the econometrician *overreacts* if $|g^* - \mu| > |\theta - \mu|$ and *underreacts* if the inequality is reversed. In other words, an estimation overreacts to information if it is further from the prior mean than the real state is.

Proposition 3.5.7 (Guess Characterization) *Let* $\boldsymbol{a}(s_i, \rho_i) = \frac{\rho_\mu \mu + \rho_i s_i}{\rho_\mu + \rho_i}$. *Under Assumption 3.5.6,* $g(a^N) \xrightarrow{a.s.} g^*$, *where*

1. *If* $\mu = \theta$, $g^* = \theta$

2. *If* $\mu \neq \theta$, *then* $\exists \tilde{\rho} < \tilde{\tilde{\rho}}$ *such that*

   - $\rho^* \leq \tilde{\rho} \Rightarrow |g^* - \mu| > |\theta - \mu|$ *and the agent underreacts to observed actions;*
   - $\rho^* \geq \tilde{\tilde{\rho}} \Rightarrow |g^* - \mu| < |\theta - \mu|$ *and the agent overreacts to observed actions;*
   - $\tilde{\rho} < \rho^* < \tilde{\tilde{\rho}} \Rightarrow$ *the agent underreacts to observed actions if* $|\theta - \mu|$ *is small and overreacts if* $|\theta - \mu|$ *is large,*

   where:

$$\tilde{\rho} = \frac{2\underline{\rho}\overline{\rho}}{\underline{\rho} + \overline{\rho}} \qquad \tilde{\tilde{\rho}} = \underline{\rho} H\left(\overline{m}(\tilde{\tilde{\rho}}, \mu)\right) + \overline{\rho}\left(1 - H\left(\overline{m}(\tilde{\tilde{\rho}}, \mu)\right)\right).$$

Proposition 3.5.7 reveals that whether the decision maker over- or underreacts depends on the true precision of the signals and possibly the realization of the state θ. Roughly speaking, the optimal robust estimate corresponds to the decision maker trying to backtrack the mean of the unobservable signals from the mean of observed actions. Because signals are unbiased, their unobservable mean is effectively θ, the state the econometrician aims to estimate. When ρ^* is high, θ is relatively close to the mean of actions, because the agents place a high weight on their unbiased signals when choosing their actions. However, the econometrician does not know the real precision, so she backtracks signals from actions using, roughly, the same method regardless of what ρ^* is. Therefore, the direction of her estimation error depends on the true precision.

Finally, as the prior precision ρ_μ changes, the acuraccy of the optimal estimate is not monotonic. The estimation error is related to how actions are contaminated by the prior, making it impossible for the agent to disentangle the effect of the prior from the effect of individual information. When the prior is extremely imprecise, $\rho_\mu \approx 0$, this contamination is minimal, and the optimal estimate is approximately equal to the one in the observable signals example: the econometrician estimates correctly. On the other hand, when the precision of the prior grows to infinity, $\rho_\mu \to \infty$, the agent

also estimates correctly by essentially disregarding the information in the observed actions. For intermediate values, however, Corollary 3.5.5 implies that the estimate is wrong almost surely. In other words, the accuracy of the estimate is not monotonic with the precision of the prior: better information ex-ante does not guarantee a more correct estimate asymptotically.

3.6 Discussion

To maintain tractability and clarity, our analysis has relied on four main assumptions: (i) The decision maker only knows the highest possible and lowest possible precisions of each information source and nothing else; (ii) both the state and signals follow normal distributions; (iii) the decision maker is an MEU maximizer regarding ambiguity; and (iv) the decision maker adopts full Bayesian updating. In this section, we briefly argue our main result — that ambiguity does not vanish asymptotically — remains valid when we relax the first three assumptions. Hence, the essential assumption is the updating rule under ambiguity.

Information about Precisions. Consider the case in which the decision maker has more information about the precisions of her information sources ex-ante. Specifically, the decision maker knows two groups of information sources exist. Group 1 consists of a fraction $\alpha \in [0, 1]$ of information sources with shared high precision $\overline{\rho}$, and Group 2 consists of fraction $1 - \alpha$ with shared low precision $\underline{\rho}$. The decision maker does not know which group a particular information source belongs to. Recall the optimization problem of the decision maker can be interpreted as a zero-sum game between her and nature. The decision maker's additional information heavily restricts nature's choices on precisions. However, when $\alpha \in (0, 1)$, nature can still induce the decision maker to have a relatively high (low) posterior mean of the state by assigning high signals to Group 1 (Group 2) subject to the new constraint. Hence, even with the additional restrictions, ambiguity will not asymptotically vanish. The asymptotic estimate of the decision maker will be correct only with observable signals and a symmetric loss

function, and incorrect otherwise.[6] Regardless of whether her estimate is correct, for any $\alpha \in (0,1)$, the decision maker faces ambiguity, and thus, suffers from a loss. Consequently, an decision maker who believes all her information sources to have minimal precision $\underline{\rho}$, is better off than an decision maker who believes a fraction of her information sources have precision $\overline{\rho} > \underline{\rho}$.

Distributions. Typically, ambiguity does not vanish even when the state and signals are not normally distributed. In particular, as long as the decision maker's belief set is rich enough, ambiguity about the state will persist asymptotically. We expand on this point next. For general distributions, the precision of each signal is no longer fully captured by the reciprocal of its variance. To extend our model to other distributions, assume the decision maker considers a set of likelihood functions for each information source. As in the main model, each allocation of likelihoods to information sources defines a belief for the agent. Under full Bayesian updating for each belief, the agent forms a posterior on the state. As long as two different beliefs result in two different posterior means, our results hold. This condition is generic: if a non-singleton set of likelihood functions does not satisfy this property, one of the functions can be perturbed so that the property holds and ambiguity does not vanish. Finally, departing from normal distributions, higher moments of the posterior no longer necessarily vanish. Thus, in addition to the induced ambiguity, the agent might face risk even asymptotically.

Ambiguity Preferences. We can extend the decision maker's preference under ambiguity. As long as the ambiguity the decision maker faces takes the form of a set of beliefs over the state and signals and she adopts the full Bayesian updating rule upon receiving signals, Theorem 3.3.5 still holds. Indeed, our analysis on asymptotic beliefs does not rely on the specific utility function of the decision maker. For instance, ambiguity does not vanish if the decision maker has an uncertainty-averse preference

[6]Note that if $\alpha = 1$ or 0, that is, if the decision maker knows that all of her information sources are precise or imprecise, asymptotic learning is successful, and the decision maker has asymptotic loss 0 because no ambiguity exists.

introduced by Cerreia-Vioglio et al. (2011), which incorporates most ambiguity-averse preferences used in the literature, including variational preferences (Klibanoff et al., 2005), smooth ambiguity preferences (Maccheroni et al., 2006), and, of course, MEU.

Updating Rule. Finally, we note our result *does* rely on the updating rule under ambiguity. An alternative to the full Bayesian updating rule is the maximum-likelihood rule. Unlike full Bayesian updating, where the decision maker applies Bayes' rule to the entire set of priors, the decision maker with the maximum-likelihood rule would discard priors that do not ascribe the maximal probability to the observed signals and update the remaining priors according to Bayes' rule. Hence, the maximum-likelihood rule suggests that ambiguity might vanish even with one single signal.

Appendix A

Appendices for Chapter 1

A.1 Asymmetric Narrow Bracketing in Infinite Horizon

In this appendix, I briefly discuss how to extend the asymmetric narrow bracketing model in (1.4) to one with multiple periods. Assume that the consumption space in each period $t = 1, ..., T$ is a compact interval $C \subset \mathbb{R}_+$, where T can be $+\infty$. The set of deterministic consumption paths is C^T with a generic element $\mathbf{c} = (c_t)_{t=1}^T$. For each consumption path $\mathbf{c} \in C^T$, I denote the subsequence of consumption in the first t periods as $\mathbf{c}^t = (c_\tau)_{\tau=1}^t$.

The preference relation is defined on the lottery space $\mathcal{P} = \mathcal{L}(C^T)$. Here I allow for lotteries with infinite supports to accommodate applications in finance. For each lottery P, denote $P_{[t]}$ as the marginal lottery in the first t periods, $1 \leq t < T$. For each subsequence of consumption $\mathbf{c}^t$ in the support of $P_{[t]}$, I define $\phi(P|\mathbf{c}^t)$ as the conditional lottery starting from period $t+1$, given that consumption in the first t periods are $\mathbf{c}^t$. When $T < +\infty$, $\phi(P|\mathbf{c}^t) \in \mathcal{L}(C^{T-t})$ and when $T = +\infty$, $\phi(P|\mathbf{c}^t) \in \mathcal{L}(C^\infty)$. Note that for each finite T, $\mathcal{L}(C^{T-t})$ is homeomorphic to a subset of $\mathcal{L}(C^\infty)$ where the consumption levels are always 0 from period $t+1$ on. So I will focus on the case with an infinite horizon.

The following notions are adapted from recursive preferences on temporal lotteries (Chew and Epstein, 1991, Bommier, Kochov, and Le Grand, 2017) to my framework. For each $V : \mathcal{P} := \mathcal{L}(C^\infty) \to \mathbb{R}$ and $p \in \mathcal{P}$, denote

$$m_V(P)(B) \equiv P_1\Big\{c \in C : V(c, \phi(P|c) \in B)\Big\}, \forall B \in \mathcal{B}(V(\mathcal{P}))$$

where $V(\mathcal{P}) \subset \mathbb{R}$ is the image of V on $\mathcal{P}$ and $\mathcal{B}(V(\mathcal{P}))$ is the set of all Borel subsets of $V(\mathcal{P})$. Then $m_V(P)$ is a probability measure over utilities conditional on the current consumption. Now I define the *recursive preference over lotteries* as $V : \mathcal{P} \to \mathbb{R}$ with

$$V(P) = I(m_V(P)),$$
$$V(c, q) = W(c, V(q)),$$

where $m_V(P)$ is defined as above, $I : \mathcal{L}(\mathbb{R}) \to \mathbb{R}$ is a *certainty equivalent*, that is, I is continuous, increasing with respect to first order stochastic dominance and $I(x) = x$ for each $x \in \mathbb{R}$, $W : C \times \mathbb{R} \to \mathbb{R}$ is continuous and strictly increasing in the second argument. Unlike the recursive preferences in Chew and Epstein (1991) and Bommier, Kochov, and Le Grand (2017), V is defined on a different domain and can be discontinuous.

In order to get the CRRA-CES functional form, I can set $I = \phi^{-1} \circ \mathbb{E} \circ \phi$ with $\phi(x) = x^{\alpha/\rho}$ and $W(c, v) = (1-\delta)c^\rho + \delta v$, where $\rho < 1, 0 \neq \alpha < 1$ and $0 < \delta < 1$. The recursive preference is equivalent to the following recursion of value functions (up to a monotonic transformation):

$$V^{ANB}(P) = \frac{1}{\alpha}\mathbb{E}_{P_1}\Big(U_1^\alpha\Big), \tag{A.1}$$
$$U_t^\rho = (1-\delta)c_t^\rho + \delta\Big[\mathbb{E}_{\phi(P|\mathbf{c}^t)_1}\Big(U_{t+1}^\alpha\Big)\Big]^{\frac{\rho}{\alpha}}, \tag{A.2}$$

where U_t is the value in period t and the expectation is computed with respect to $\phi(P|\mathbf{c}^t)_1$, which is the probability distribution of consumption levels in period $t+1$ conditional on the consumption path in the first t periods $\mathbf{c}^t$.

Then I explore the implications of asymmetric narrow bracketing in a standard asset pricing problem. The consumer is endowed with initial wealth $W_1 > 0$ in period 1 and chooses the consumption level and saving level in each period. Let S denote the finite state space in each period $t \geq 2$ and $\Omega = S^\infty$ denote the space of state sequences. The consumer has a prior belief over Ω. For each $s^\infty \in \Omega$, I denote by s^t the history of states from period 2 to period t for each $t \geq 2$. Let S^t be the set of all histories till period t.

The consumer's preference is represented by V^{ANB} in (A.1). She chooses a *consumption plan* $\{c_t\}_{t\geq 1}$ to maximizes her utility, where $c_1 \in X$ and $c_t : S^t \to X$ for each $t \geq 2$. Given the history of states s^t, the gross return on wealth from period $t-1$ to period t is $R_{w,t}(s^t) > 0$. Then the wealth dynamics are represented by the following equation:

$$W_{t+1}(s^{t+1}) = R_{w,t+1}(s^{t+1})(W_t(s^t) - c_t(s^t)).$$

I assume that the wealth level always lies in C. I say that a consumption plan $\{c_t\}_{t\geq 1}$ is feasible given initial wealth W_1 if if $c_t(s^t) \leq W_t(s^t)$ for all s^t, t. Each feasible consumption plan $\{c_t\}_{t\geq 1}$ induces a lottery $P \in \mathcal{P}$ where the consumption in the first period is deterministic. Using V^{ANB} in (A.1), I can define a utility function over feasible consumption plans as $\hat{V}^{ANB}(\{c_t\}_{t\geq 1})$. Then the optimization problem of the consumer is

$$J^{ANB}(W_1) = \sup\left\{\hat{V}^{ANB}(\{c_t\}_{t\geq 1}) : \{c_t\}_{t\geq 1} \text{ is feasible given } W_1\right\}.$$

To facilitate the comparison of ANB and EZ, I can similarly consider a consumer with a CRRA-CES EZ recursive utility function $\hat{V}^{EZ}$ over feasible consumption plans and the optimal value $J^{EZ}(W_1)$. I assume that RRA>1/EIS, then the EZ consumer has a preference over early resolution of uncertainty, while the ANB consumer exhibits indifference to temporal resolution of uncertainty. The following proposition shows that the two utility functions lead to the same optimal value.

Proposition A.1.1 *Assume RRA > 1/EIS, i.e., $\rho > \alpha$. Then for each $W_1 > 0$, $J^{ANB}(W_1) = J^{EZ}(W_1)$ exists and there exist consumption plans $\{c_t\}_{t\geq 1}$ and $\{c_t^n\}_{t\geq 1, n\geq 1}$ feasible given W_1 such that $c_1^n \to c_1, c_t^n(s^t) \to c_t(s^t)$ as n goes to infinity for each $t \geq 2$, and*

(i). $J^{ANB}(W_1) = \lim_{n\to\infty} \hat{V}^{ANB}(\{c_t^n\}_{t\geq 1}) = \hat{V}^{EZ}(\{c_t\}_{t\geq 1}) = J^{EZ}(W_1)$;

(ii). For each $n \geq 1, t \geq 2$, c_t^n is injective on Ω_t, that is, $c_t^n(s^t) \neq c_t^n(\hat{s}^t)$ if $s^t \neq \hat{s}^t$.

Here I sketch the proof of Proposition A.1.1. Recall that in Section 1.7.1, I observe that, in the setup with two periods, V^{ANB} agrees with V^{EZ} when risk does not resolve early. This observation remains valid with an infinite horizon. On the one hand, since the EZ consumer has a preference for early resolution of risk, I know that $\hat{V}^{ANB}(\{c_t\}_{t\geq 1}) \leq \hat{V}^{EZ}(\{c_t\}_{t\geq 1})$ for each feasible consumption plan $\{c_t\}_{t\geq 1}$ and $J^{ANB}(W_1) \leq J^{EZ}(W_1)$. On the other hand, for a consumption plan $\{c_t\}_{t\geq 1}$ where c_t is injective on Ω_t for each t, the consumption history contains the same information as state history and there is no early resolution of risk. In this case, $\hat{V}^{ANB}(\{c_t\}_{t\geq 1}) = \hat{V}^{EZ}(\{c_t\}_{t\geq 1})$. By continuity of V^{EZ} and hence $\hat{V}^{EZ}$, I can find a sequence of consumption plans $\{c_t^n\}_{t\geq 1, n\geq 1}$ with injective consumption functions to approximate the optimal value of the EZ consumer. This implies $J^{ANB}(W_1) = J^{EZ}(W_1)$.

A directly corollary of Proposition A.1.1 is that my model of asymmetric narrow bracketing has the same implications in asset pricing as the EZ model when RRA>1/EIS, which is the common parametric assumption in most applications of EZ preferences in finance and macroeconomics. Following the asset pricing literature, let $\gamma := 1 - \alpha$ be the RRA, $\psi := \frac{1}{1-\rho}$ be the EIS and $\theta := \frac{\alpha}{\rho}$.

Corollary A.1.2 *Assume RRA>1/EIS, i.e., $\rho > \alpha$. Denote $\mathbb{E}_t$ as the expectation with respect to s^t. Then for each $W_1 > 0$,*

(i). Euler equation:

$$\lim_{n\to\infty} \mathbb{E}_t\Big[\delta^\theta (\frac{c_{t+1}^n}{c_t^n})^{-\frac{\theta}{\psi}} R_{w,t+1}^\theta\Big] = \mathbb{E}_t\Big[\delta^\theta (\frac{c_{t+1}}{c_t})^{-\frac{\theta}{\psi}} R_{w,t+1}^\theta\Big] = 1;$$

(ii). Asset pricing formula: for each asset i with gross return $R_{i,t}$,

$$\lim_{n\to\infty} \mathbb{E}_t\Big[\delta^\theta (\frac{c_{t+1}^n}{c_t^n})^{-\frac{\theta}{\psi}} R_{w,t+1}^{-(1-\theta)} R_{i,t+1}\Big] = \mathbb{E}_t\Big[\delta^\theta (\frac{c_{t+1}}{c_t})^{-\frac{\theta}{\psi}} R_{w,t+1}^{-(1-\theta)} R_{i,t+1}\Big] = 1.$$

I end this section with a brief discussion about the case with RRA¡1/EIS, i.e., when the EZ consumer prefers late solution of risk. In this case, $J^{ANB}(W_1)$ might be strictly higher than $J^{EZ}(W_1)$ due to the discontinuity of V^{ANB} on $\mathcal{P}$, which invalidates my proof of Proposition A.1.1. Interestingly, the ANB consumer might have excessive demand for consumption smoothing across different states of the world in the same period. However, the derivation of the optimal value and Euler equation is much less tractable and I leave a formal treatment of this case for future study.

A.2 Proof of Theorem 1.4.1

Necessity. V^{CN}, V^{ANB-CN} and V^{FNB} trivially satisfy Weak Order and Correlation Neutrality. Given Correlation Neutrality, Continuity is equivalent to Continuity 3, which is guaranteed by continuity of w, v_1 and v_2 in the three utility functions. Similarly, Monotonicity of $\succsim$ is guaranteed by monotonicity of w, v_1 and v_2.

Now I verify Independence*. If $\succsim$ is represented by V^{CN}, then for any $P, Q, R, S \in \hat{\mathcal{P}}$ with $P_i = R_i, Q_j = S_j$ for some $i, j \in \{1, 2\}$ and $\alpha \in (0, 1)$,

$$P \succsim Q, R \sim S \Longleftrightarrow \alpha P + (1-\alpha)R \succsim \alpha Q + (1-\alpha)R.$$

Hence $\succsim$ satisfies Independence*.

If $\succsim$ is represented by V^{FNB}, then $\succsim$ satisfies part (i) of Independence* since $\succsim_{1|p}$ and $\succsim_{2|q}$ admit EU representations for any $p \in \mathcal{L}^0(X_2), q \in \mathcal{L}^0(X_1)$. For part (ii) and (iii), $P_i = R_i, P_{-i} \sim_{-i|x} R_{-i}$ for some $i \in \{1, 2\}$ and $x \in X_i$ implies $P \sim R \sim \alpha P + (1-\alpha)R$ for any $\alpha \in (0, 1)$. Similarly, $Q \sim S \sim \alpha Q + (1-\alpha)S$ for any $\alpha \in (0, 1)$. Hence $P \succ Q$ suggests that $\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$.

If $\succsim$ is represented by V^{ANB}, then $\succsim_{2|q}$ is represented by EU v_2 for any $q \in \mathcal{L}^0(X_1)$ and $P_1 = R_1, P_2 \sim_{2|x} R_2$ for some $x \in X_1$ implies $P \sim R$. Hence it suffices to show

that $\succsim$ satisfies part (i) and part (ii) of Independence* when $P_2 = R_2, Q_2 = S_2$, which is guaranteed since for any $P, Q, R, S \in \hat{\mathcal{P}}$ with $P_2 = R_2, Q_2 = S_2$ and $\alpha \in (0,1)$,

$$P \succsim Q, R \sim S \Longleftrightarrow \alpha P + (1-\alpha)R \succsim \alpha Q + (1-\alpha)R.$$

Sufficiency. Throughout the proof, for any set of lotteries $\mathcal{A} \subseteq \mathcal{P}$, denote $\max_{\succsim} \mathcal{A} = \{P \in \mathcal{A} : P \succsim Q, \forall Q \in \mathcal{A}\}$ whenever it is well-defined. $\max_{\succsim} \mathcal{A}$ is the set of most preferred lotteries in $\mathcal{A}$ under $\succsim$. Finally, for any$A \subseteq \mathbb{R}$, denote A^o as its interior.

Assume all axioms hold. First, by Correlation Neutrality, $P \sim (P_1, P_2)$ for all $P \in \mathcal{P}$ and it suffices to focus on the restriction of $\succsim$ on product lotteries $\hat{\mathcal{P}}$. Then Continuity implies that $\succsim$ satisfies strong continuity, that is, for all $Q \in \mathcal{P}$, the sets $\{P \in \mathcal{P} : P \succ Q\}$ and $\{P \in \mathcal{P} : Q \succ P\}$ are open subsets of $\mathcal{P}$. The following lemma guarantees that the conditional preference $\succsim_{i|q}$ admits an EU representation for each $i = 1, 2$ and $q \in \mathcal{L}^0(X_{-i})$.

Lemma A.2.1 *For each $i = 1, 2$ and $q \in \mathcal{L}^0(X_{-i})$, $\succsim_{i|q}$ admits an EU representation with index $v_{i|q}$, which is continuous and unique up to a positive affine transformation. Moreover, if $q \in X_{-i}$, then $v_{i|q}$ can be chosen to be strictly monotonic.*

Proof of Lemma A.2.1 . Fix $i = 1, 2$ and $q \in \mathcal{L}^0(X_{-i})$. By Continuity and Independence*, $\succsim_{i|q}$ admits an EU representation with a continuous utility index $v_{i|q}$ defined on X_i, which is unique up to a positive affine transformation. Moreover, if $q \in X_{-i}$, that is, $q = y$ for some $y \in X_{-i}$, then by Monotonicity, $v_{i|q}$ must be strictly monotonic. ■

A direct corollary of Lemma A.2.1 guarantees the existence of "certainty equivalents".

Corollary A.2.2 *For each $P \in \hat{\mathcal{P}}$, there exists $x_1, y_1 \in X_1, x_2, y_2 \in X_2$ such that $P \sim (P_1, x_2) \sim (x_1, P_2) \sim (y_1, y_2)$.*

Proof of Corollary A.2.2 . I will prove the result for $P_1 \notin X_1, P_2 \notin X_2$. The proof for other cases is easier. By Lemma A.2.1, I can find $x_2, a, a' \in X_2$ such that either $\sum_x v_{2|P_1}(x)P_2(x) = v_{2|P_1}(x_2)$ and hence $P \sim (P_1, x_2)$ or $v_{2|P_1}(a) > \sum_x v_{2|P_1}(x)P_2(x) >$

$v_{2|P_1}(a')$. In the latter case, since $v_{2|P_1}$ is continuous and X_2 is a closed interval, there exists $x_2 \in X_2$ where $v_{2|P_1}(x_2) = \sum_x v_{2|P_1}(x)P_2(x)$, which implies $P \sim (P_1, x_2)$. Similarly, I can find $x_1 \in X_1$ with $P \sim (x_1, P_2)$. Now let $y_2 = x_2$. Repeat the above arguments for product lottery (P_1, x_2) and I can find $y_1 \in X_1$ with $(y_1, y_2) \sim (P_1, x_2) \sim P$. ■

The next lemma shows that under Continuity, Independence* can be strengthened.

Lemma A.2.3 *(i). For each $P, Q, R \in \hat{\mathcal{P}}$ with $P \succsim R \succsim Q$, $P \succ Q$ and $P_i = Q_i$ for some $i \in \{1, 2\}$, then there exists a unique $\lambda \in [0, 1]$ such that $R \sim \lambda P + (1 - \lambda)Q$. (ii). For $P, Q, R, S \in \hat{\mathcal{P}}$, $\alpha \in (0, 1)$, $i, j \in \{1, 2\}$ and $x \in X_i, y \in X_j$, if $P_i = R_i, Q_j = S_j, P_{-i} \sim_{-i|x} R_{-i}$ and $Q_{-j} \sim_{-j|y} S_{-j}$, then*

$$P \sim Q, R \sim S \Longrightarrow \alpha P + (1 - \alpha)R \sim \alpha Q + (1 - \alpha)S,$$
$$P \succ Q, R \succ S \Longrightarrow \alpha P + (1 - \alpha)R \succ \alpha Q + (1 - \alpha)S.$$

(iii). For $P, Q \in \hat{\mathcal{P}}$ and $x \in X_1$, if $R = (P_1, Q_2)$ or (Q_1, P_2) and $P_2 \sim_{2|x} Q_2$, then

$$P \sim Q \Longrightarrow \alpha P + (1 - \alpha)R \sim \alpha Q + (1 - \alpha)R.$$

Proof of Lemma A.2.3. (i). Denote $I = \{\eta \in [0, 1] : R \succ \eta P + (1 - \eta)Q\}$ and $\lambda = \sup I$. λ is well-defined as I is bounded. I claim that $\lambda P + (1 - \lambda)Q \sim R$. If $\lambda P + (1 - \lambda)Q \succ R$, then $\lambda > 0$ and by Continuity, there exists $\varepsilon > 0$ with $(\lambda - \varepsilon)P + (1 - \lambda + \varepsilon)Q \succ R$. This implies $\lambda - \varepsilon \notin I$. Since $P \succ Q$ and $P_i = Q_i$ for some i, by Independence*, for any $\alpha, \beta \in [0, 1]$, $\alpha P + (1 - \alpha)Q \succ \beta P + (1 - \beta)Q$ if any only if $\alpha > \beta$, which implies $[\lambda - \varepsilon, \lambda] \cap I = \emptyset$ and leads to a contradiction with $\lambda = \sup I$. If instead $R \succ \lambda P + (1 - \lambda)Q$, then there exists $\varepsilon > 0$ with $R \succ (\lambda + \varepsilon)P + (1 - \lambda - \varepsilon)Q$ and hence $\lambda + \varepsilon \in I$, which again contradicts with the definition of λ. Uniqueness of λ is guaranteed by its definition.

(ii). Suppose $P \sim Q, R \sim S$. If $P \sim R$, then the the result trivially holds as $\alpha P + (1 - \alpha)R \sim R \sim Q \sim \alpha Q + (1 - \alpha)S$. Without loss of generality, assume $P \succ R$. Then $Q \succ S$ and $P \succ \alpha P + (1 - \alpha)R \succ R$, $Q \succ \alpha Q + (1 - \alpha)S \succ S \sim R$ for any

$\alpha \in (0,1)$. Suppose by contradiction that $\alpha P+(1-\alpha)R \succ \alpha Q+(1-\alpha)S$. By part (i), there exists a unique $\lambda \in (0,1)$ with $\alpha Q+(1-\alpha)S \sim \lambda(\alpha P+(1-\alpha)R)+(1-\lambda)R = \alpha\lambda P+(1-\alpha\lambda)R$. Notice that $Q \sim P \succ \lambda P+(1-\lambda)R$, $S \sim R$, $Q_j = S_j$, $Q_{-j} \sim_{-j|y} S_{-j}$, $(\lambda P + (1-\lambda)R)_i = P_i = R_i$ and $(\lambda P + (1-\lambda)R)_{-i} \sim_{-i|x} R_{-i}$. The last property follows from Lemma A.2.1. Then Independence* implies

$$\alpha Q + (1-\alpha)S \succ \alpha(\lambda P + (1-\lambda)R) + (1-\alpha)R = \alpha\lambda P + (1-\alpha\lambda)R,$$

which leads to a contradiction. The case for $\alpha P + (1-\alpha)R \prec \alpha Q + (1-\alpha)S$ is symmetric.

Now assume $P \succ Q, R \succ S$. If $P \sim R$, then the result holds as $\alpha P + (1-\alpha)R \sim P \succ \max_{\succsim}\{Q,S\} \succsim \alpha Q + (1-\alpha)S$ for all $\alpha \in (0,1)$. Without loss of generality, suppose $P \succ R$. If $R \succsim Q$, then $\alpha P+(1-\alpha)R \succ R \succsim \max_{\succsim}\{Q,S\} \succsim \alpha Q+(1-\alpha)S$ for all $\alpha \in (0,1)$. If instead $Q \succ R$, then $P \succ Q \succ R \succ S$. By part (i), I can find $\lambda \in (0,1)$ such that $R \sim \lambda Q + (1-\lambda)S := S' \succ S$. Since $S'_j = S_j = Q_j$ and $S'_{-j} = \lambda Q_{-j} + (1-\lambda)S_{-j} \sim_{-j} Q_{-j}$, Independence* implies for any $\alpha \in (0,1)$,

$$\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S' \succ \alpha Q + (1-\alpha)S.$$

(iii). The proof is the same as the first part of the proof for (ii) and omitted. ■

A tuple $(P,Q,R,S) \in \hat{\mathcal{P}}^4$ is called *proper* if $P_i = R_i$, $Q_j = S_j$ for some $i,j \in \{1,2\}$ and $P \succsim R, Q \succsim S$. A proper tuple (P,Q,R,S) satisfies the *independence property* if one of the following conditions holds:

- $P \succ Q, R \sim S$ and for all $\alpha \in (0,1)$, $\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$;
- $P \sim Q, R \succ S$ and for all $\alpha \in (0,1)$, $\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$;
- $P \sim Q, R \sim S$ and for all $\alpha \in (0,1)$, $\alpha P + (1-\alpha)R \sim \alpha Q + (1-\alpha)S$;
- $P \succ Q, R \succ S$ and for all $\alpha \in (0,1)$, $\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$.

Now I introduce additional notations that will be useful later on. For each $x_i \in X_i$, $i = 1,2$, denote $\Pi^i(x_i|x_{-i})$ as the set of marginal lotteries in source i with certainty

equivalent x_i given background outcome x_{-i} in source $-i$. Formally, $\Pi^i(x_i|x_{-i}) = \{p \in \mathcal{L}^0(X_i) : p \sim_{i|x_{-i}} x_i\}$. For any $P, Q \in \hat{\mathcal{P}}$ with $P \succsim Q$, let $[Q, P]$ denote the set of product lotteries whose utilities lie between P and Q, that is, $[Q, P] = \{S \in \hat{\mathcal{P}} : P \succsim S \succsim Q\}$. Then, for each $q_i \in \mathcal{L}^0(X_i)$ and $x_i \in X_i$, $i = 1, 2$, define

$$\Gamma_{1,q_1}(x_2|x_1) = \bigcup_{\substack{P,Q \in \{q_1\} \times \Pi^2(x_2|x_1), \\ P \succsim Q}} [Q, P] \ , \quad \Gamma_{2,q_2}(x_1|x_2) = \bigcup_{\substack{P,Q \in \Pi^1(x_1|x_2) \times \{q_2\}, \\ P \succsim Q}} [Q, P].$$

Intuitively, $\Gamma_{1,q_1}(x_2|x_1)$ includes all product lotteries whose utilities are bounded by lotteries in $\{q_1\} \times \Pi^2(x_2|x_1)$. The following lemma states that any lottery in $\Gamma_{1,q_1}(x_2|x_1)$ is indifferent to some lottery in $\{q_1\} \times \Pi^2(x_2|x_1)$.

Lemma A.2.4 *Fix $q_1 \in \mathcal{L}^0(X_1), q_2 \in \mathcal{L}^0(X_2)$ and $x_1 \in X_1, x_2 \in X_2$.*
(i). For each $P \in \Gamma_{1,q_1}(x_2|x_1)$, there exists $P' \in \{q_1\} \times \Pi^2(x_2|x_1)$ with $P' \sim P$;
(ii). For each $P \in \Gamma_{2,q_2}(x_1|x_2)$, there exists $P' \in \Pi^1(x_1|x_2) \times \{q_2\}$ with $P' \sim P$.

Proof of Lemma A.2.4. (i). By definition, there exists $Q, Q' \in \{q_1\} \times \Pi^2(x_2|x_1)$ with $Q \succsim P \succsim Q'$. As $Q_1 = Q'_1 = q_1$, by Lemma A.2.3, there exists $\lambda \in [0, 1]$ with $P \sim \lambda Q + (1 - \lambda)Q'$. By Lemma A.2.1, $(\lambda Q + (1 - \lambda)Q')_2 \sim_{2|x_1} Q_2 \sim_{2|x_1} Q'_2$. Denote $P' = \lambda Q + (1 - \lambda)Q'$. Hence $P' \sim P$, $P' \in \{q_1\} \times \Pi^2(x_2|x_1)$. The proof for (ii) is similar. ■

The rest of proof can be decomposed into several cases.

Case 1. There exists $(z_1, z_2) \in X_1 \times X_2$ such that $(p, q) \sim (x, y)$ for all $(x, y) \in X_1 \times X_2$ and $(p, q) \in \Pi^1(x|z_2) \times \Pi^2(y|z_1)$. The following lemma shows that $\succsim$ must admit a FNB representation.

Lemma A.2.5 *In Case 1, $\succsim$ is represented by V^{FNB}.*

Proof of Lemma A.2.5. By assumption, the conditional preferences $\succsim_{1|p}$ and $\succsim_{2|q}$ are independent of $p \in \mathcal{L}^0(X_2)$ and $q \in \mathcal{L}^0(X_1)$. I denote by $\succsim_1$ and $\succsim_2$ the two conditional preferences respectively. By Lemma A.2.1, for $i = 1, 2$, I can find a continuous and strictly monotonic v_i as the EU index of $\succsim_i$. Since X_i is a closed

interval, the certainty equivalent function CE_{v_i} is well-defined. For any $(p,q) \in \hat{\mathcal{P}}$, I know $(p,q) \sim (CE_{v_1}(p), CE_{v_2}(q))$. I define $\hat{\succsim}$ as the restriction of $\succsim$ on $X_1 \times X_2$. By Continuity, $\hat{\succsim}$ is continuous. Then Debreu's Theorem implies that $\hat{\succsim}$ is represented a continuous utility function w. Monotonicity guarantees that w is strictly monotonic. Therefore w, v_1 and v_2 are continuous and strictly monotonic and for all $P, Q \in \mathcal{P}$,

$$\begin{aligned} P \succsim Q &\iff (CE_{v_1}(P_1), CE_{v_2}(P_2)) \hat{\succsim} (CE_{v_1}(Q_1), CE_{v_2}(Q_2)) \\ &\iff w(CE_{v_1}(P_1), CE_{v_2}(P_2)) \geq w(CE_{v_1}(Q_1), CE_{v_2}(Q_2)) \\ &\iff V^{FNB}(P) \geq V^{FNB}(Q). \end{aligned}$$

■

Case 2: There exists $(z_1, z_2) \in X_1 \times X_2$ such that $(p,q) \sim (p,q')$ for all $(x,y) \in X_1 \times X_2$, $p \in \Pi^1(x|z_2)$ and $q, q' \in \Pi^2(y|z_1)$, and $\exists\ (x_0, y_0) \in X_1 \times X_2$, $p_0, p_0' \in \Pi^1(x_0|z_2)$ and $q_0 \in \Pi^2(y_0|z_1)$ with $(p_0, q_0) \succ (p_0', q_0)$.

By assumption, $\succsim_{2|q}$ is independent of $q \in \mathcal{L}^0(X_1)$ and I can denote it as $\succsim_2$. By Lemma A.2.1, $\succsim_2$ admits a continuous and strictly monotonic EU index v_2. Then for any $(p,q), (p',q') \in \hat{\mathcal{P}}$ with $q \sim_2 q'$, $(p,q) \succsim (p',q')$ if and only if $(p, CE_{v_2}(q)) \succsim (p', CE_{v_2}(q))$. Hence I can focus on the restriction of $\succsim$ on $\mathcal{L}^0(X_1) \times X_2$.

Lemma A.2.6 *In Case 2, there exists $z_0, z_0' \in X_2$ such that for any $x \in X_1^o$, I can find $p_x, p_x' \in \Pi^1(x|z_0')$ with $(p_x, z_0) \succ (p_x', z_0)$ and $q_x, q_x' \in \Pi^1(x|z_0)$ with $(q_x, z_0') \succ (q_x', z_0')$.*

Proof of Lemma A.2.6. By assumption, there exist $z_0, z_0' \in X_2$, $x_0 \in X_1$ and $p_0, p_0' \in \Pi^1(x_0|z_0')$ such that $(p_0, z_0) \succ (p_0', z_0)$. Suppose by contradiction that there exists $x_1 \in X_1^o$ with $(p_1, z_0) \sim (p_1', z_0)$ for all $p_1 \sim_{1|z_0'} p_1' \sim_{1|z_0'} x_1$. By Lemma A.2.1, $\succsim_{1|z_0}$ and $\succsim_{1|z_0'}$ are represented by EU $v_{1|z_0}$ and $v_{1|z_0'}$ respectively. Since $x_1 \in X_1^o$, there exists $\overline{x}, \underline{x} \in X_i^o$ with $\overline{x} > x_1 > \underline{x}$. As $v_{1|z_0}$ and $v_{1|z_0'}$ are unique up to a positive affine transformation, I can set $v_{1|z}(\underline{x}) = v_{1|z'}(\underline{x})$ and $v_{1|z_0}(x_1) = v_{1|z_0'}(x_1)$. For any $x \in X_1$ with $x > x_1$, there exists $\alpha \in (0,1)$ with $\alpha\delta_x + (1-\alpha)\delta_{\underline{x}} \sim_{1|z_0'} x_1$. Then $\alpha v_{1|z_0}(x) + (1-\alpha) v_{1|z_0}(\underline{x}) = v_{1|z_0}(x_1) = v_{1|z_0'}(x_1) = \alpha v_{1|z_0'}(x) + (1-\alpha) v_{1|z_0'}(\underline{x})$. Since

$v_{1|z_0}(\underline{x}) = v_{1|z_0'}(\underline{x})$ and $\alpha \in (0,1)$, $v_{1|z_0}(x) = v_{1|z_0'}(x)$ for all $x > x_1$. Specifically, I have $v_{1|z_0}(\overline{x}) = v_{1|z_0'}(\overline{x})$. Repeat the above arguments and I can show $v_{1|z_0}(x) = v_{1|z_0'}(x)$ for all $x < x_1$. Thus $v_{1|z_0} \equiv v_{1|z_0'}$, contradicting with $(p_0, z_0) \succ (p_0', z_0)$ and $(p_0, z_0') \sim (p_0', z_0')$. Hence for any $x \in X_1^o$, I can find $p_x, p_x' \in \Pi^1(x|z_0')$ with $(p_x, z_0) \succ (p_x', z_0)$. The existence of q_x, q_x' can be shown similarly. ■

For each $z \in X_2$, denote

$$\Sigma^2(z) := \{y \in X_2 : \exists\ x \in X_1 \text{ and } p, p' \in \Pi^1(x|z) \text{ s.t. } (p, y) \succ (p', y)\}.$$

To interpret $\Sigma^2(z)$, note that $y \in \Sigma^2(z)$ if and only if there exists $q \sim_{2|z} y$ such that $\succsim_{1|q}$ is not identical to $\succsim_{1|z}$. The following lemma follows.

Lemma A.2.7 *In Case 2, (i) $\Sigma^2(z)$ is open in the relative topology of X_2 for each $z \in X_2$; (ii) there exists $z_0, z_0' \in X_2$ such that $\Sigma^2(z_0) \cup \Sigma^2(z_0') = X_2$ and $\Sigma^2(z_0) \cap \Sigma^2(z_0') \neq \emptyset$.*

Proof of Lemma A.2.7. (i). For $z \in X_2$, if $\Sigma^2(z) \neq \emptyset$, then $\Sigma^2(z)$ is open in the relative topology of X_2 due to Continuity.

(ii). By Lemma A.2.6, there exists $z_0, z_0' \in X_2$ with $\succsim_{1|z_0} \neq \succsim_{1|z_0'}$. This implies $z_0 \in \Sigma^2(z_0')$ and $z_0' \in \Sigma^2(z_0)$. For any $z \in X_2$, either $\succsim_{1|z} \neq \succsim_{1|z_0}$ or $\succsim_{1|z} \neq \succsim_{1|z_0'}$, which suggests either $z \in \Sigma^2(z_0)$ or $z \in \Sigma^2(z_0')$. Hence $\Sigma^2(z_0) \cup \Sigma^2(z_0') = X_2$. Since X_2 is a closed interval on $\mathbb{R}$, X_2 is connected and cannot be represented as the union of two disjoint non-empty open subsets. As a result, $\Sigma^2(z_0) \cap \Sigma^2(z_0') \neq \emptyset$. ■

Using the same argument in Lemma A.2.6, I have the following corollary.

Corollary A.2.8 *In Case 2, if $z' \in \Sigma^2(z)$ where $z, z' \in X_2$, then for any $x \in X_1^o$, there exists $p_x, p_x' \in \Pi^1(x|z)$ with $(p_x, z') \succ (p_x', z')$.*

My goal is to show that a proper tuple of product lotteries (P, Q, R, S) satisfies the independence property whenever $P_2 = R_2$ and $Q_2 = S_2$. The following lemma assures that I can focus on the case where $P \sim Q$, $R \sim S$.

Lemma A.2.9 *Suppose that $(P, Q, R, S) \in \hat{\mathcal{P}}^4$ is a proper tuple with $P_i = R_i, Q_j = S_j$ for some $i, j \in \{1, 2\}$. If the independence property holds for any such (P, Q, R, S)*

with $P \sim Q, R \sim S$, then the independence property holds for any such (P, Q, R, S) with $P \succsim Q, R \succsim S$.

Proof of Lemma A.2.9. Like Lemma A.2.3, it suffices to consider the case where $P \succsim Q \succ R \succsim S$. By Lemma A.2.1, there exist $\alpha \in (0, 1]$ and $\beta \in [0, 1)$ such that $P' = \alpha P + (1-\alpha)R \sim Q$ and $S' = \beta Q + (1-\beta)S \sim R$. Then the independence property holds for (P', Q, R, S'), that is, for any $\lambda \in (0, 1)$, $\lambda P' + (1-\lambda)R \sim \lambda Q + (1-\lambda)S'$. By Lemma A.2.1 and $P \succsim P'$, $S' \succsim S$, I have

$$\lambda P + (1-\lambda)R \succsim \lambda P' + (1-\lambda)R \sim \lambda Q + (1-\lambda)S' \succsim \lambda Q + (1-\lambda)S.$$

At least one of the weak preference rankings would be strict if $P \succ Q$ or $R \succ S$. ■

The next lemma shows that if the independence property holds on two sets of product lotteries respectively, then it also holds on their union.

Lemma A.2.10 *In Case 2, $(P, Q, R, S) \in (\mathcal{L}^0(X_1) \times X_2)^4$ is a proper tuple where $P \sim Q, R \sim S$, $P_2 = R_2 = y_1, Q_2 = S_2 = y_2$ with $y_1, y_2 \in X_2$. Fix any $T^i \in \hat{\mathcal{P}}$ for $i = 1, ..., 4$ with $T^4 \succ T^2 \succ T^3 \succ T^1$. If the independence property holds for any such (P, Q, R, S) with $\{P, Q, R, S\} \subseteq [T^1, T^2]$ or $\{P, Q, R, S\} \subseteq [T^3, T^4]$, then it also holds for any such (P, Q, R, S) with $\{P, Q, R, S\} \subseteq [T^1, T^4]$.*

Proof of Lemma A.2.10. Without loss of generality, I can assume $P \succ R$ and $T^1 \succsim (x, y_i), i = 1, 2$ for some $x \in X_1$, otherwise either the lemma is trivial or I can redefine T^1 without any loss. Moreover, it suffices to focus on the case where $P \sim Q \succ T^2$ and $R \sim S \prec T^3$. Fix any W_1, W_2 with $T^2 \succ W_2 \succ W_1 \succ T^3$. Then I have

$$T^4 \succsim P \sim Q \succ T^2 \succ W_2 \succ W_1 \succ T^3 \succ R \sim S \succsim T^1.$$

By Lemma A.2.1, I can find $\hat{P}, \hat{Q}, \hat{R}, \hat{S}$ such that $\hat{P} \sim \hat{Q} \sim W_2$, $\hat{R} \sim \hat{S} \sim W_1$ and $\hat{P}_2 = \hat{R}_2 = P_2 = y_1$, $\hat{Q}_2 = \hat{S}_2 = Q_2 = y_2$. Notice that $P, Q, \hat{R}, \hat{S} \in [T^3, T^4]$, where the independence property holds. Then there exists $\lambda \in (0, 1)$ such that $\lambda P + (1-\lambda)\hat{R} \sim \hat{P} \sim \lambda Q + (1-\lambda)\hat{S}$. Similarly, I can find $\lambda' \in (0, 1)$ with $\lambda'\hat{P} + (1-\lambda')R \sim \hat{R} \sim \lambda'\hat{Q} + (1-\lambda')S$.

Actually in the construction of $\hat{R}$ and $\hat{S}$, there exists $\eta_1, \eta_2 \in (0,1)$ with

$$\eta_1 P + (1-\eta_1)R \sim \hat{R} \sim \hat{S} \sim \eta_2 Q + (1-\eta_2)S.$$

I claim that $\eta_1 = \eta_2$. To see this, as $P_2 = R_2 = \hat{P}_2 = \hat{R}_2$, I know

$$\lambda' \hat{P} + (1-\lambda')R \sim \lambda\lambda' P + (1-\lambda)\lambda' \hat{R} + (1-\lambda')R \sim \hat{R}$$

which implies

$$\hat{R} \sim \frac{\lambda\lambda'}{\lambda\lambda' + (1-\lambda')} P + \frac{1-\lambda'}{\lambda\lambda' + (1-\lambda')} R$$

and hence $\eta_1 = \frac{\lambda\lambda'}{\lambda\lambda' + (1-\lambda')}$. Similarly I can show that $\eta_2 = \frac{\lambda\lambda'}{\lambda\lambda' + (1-\lambda')} = \eta_1 := \eta^{w_1}$.

A symmetric argument shows that there exists η^{w_2} with $\eta^{w_1} < \eta^{w_2} < 1$ and

$$\eta^{w_2} P + (1-\eta^{w_2})R \sim \hat{P} \sim \hat{Q} \sim \eta^{w_2} Q + (1-\eta^{w_2})S.$$

Now I consider η with $\eta^{w_1} < \eta < \eta^{w_2}$. Notice that

$$\begin{aligned} \eta P + (1-\eta)R &= \frac{\eta - \eta^{w_1}}{\eta^{w_2} - \eta^{w_1}} [\eta^{w_2} P + (1-\eta^{w_2})R] + \frac{\eta^{w_2} - \eta}{\eta^{w_2} - \eta^{w_1}} [\eta^{w_1} P + (1-\eta^{w_1})R] \\ &\sim \frac{\eta - \eta^{w_1}}{\eta^{w_2} - \eta^{w_1}} \hat{P} + \frac{\eta^{w_2} - \eta}{\eta^{w_2} - \eta^{w_1}} \hat{R}. \end{aligned}$$

Similarly,

$$\eta Q + (1-\eta)S \sim \frac{\eta - \eta^{w_1}}{\eta^{w_2} - \eta^{w_1}} \hat{Q} + \frac{\eta^{w_2} - \eta}{\eta^{w_2} - \eta^{w_1}} \hat{S}.$$

As $\hat{P} \sim \hat{Q}, \hat{R} \sim \hat{S} \in [T^3, T^4]$ and $\hat{P}_2 = \hat{R}_2$, $\hat{Q}_2 = \hat{S}_2$, the independence property holds for $(\hat{P}, \hat{Q}, \hat{R}, \hat{S})$ and hence

$$\eta P + (1-\eta)R \sim \frac{\eta - \eta^{w_1}}{\eta^{w_2} - \eta^{w_1}} \hat{P} + \frac{\eta^{w_2} - \eta}{\eta^{w_2} - \eta^{w_1}} \hat{R} \sim \frac{\eta - \eta^{w_1}}{\eta^{w_2} - \eta^{w_1}} \hat{Q} + \frac{\eta^{w_2} - \eta}{\eta^{w_2} - \eta^{w_1}} \hat{S} \sim \eta Q + (1-\eta)S.$$

Then I check the independence property for $\eta^{w_2} < \eta < 1$.

$$\hat{P} \sim \eta^{w_2} P + (1-\eta^{w_2})R = \frac{\eta^{w_2} - \eta^{w_1}}{\eta - \eta^{w_1}} [\eta P + (1-\eta)R] + \frac{\eta - \eta^{w_2}}{\eta - \eta^{w_1}} [\eta^{w_1} P + (1-\eta^{w_1})R]$$

$$\sim \frac{\eta^{w_2} - \eta^{w_1}}{\eta - \eta^{w_1}}[\eta P + (1-\eta)R] + \frac{\eta - \eta^{w_2}}{\eta - \eta^{w_1}}\hat{R}.$$

Similarly,

$$\hat{Q} \sim \eta^{w_2} Q + (1-\eta^{w_2})S \sim \frac{\eta^{w_2} - \eta^{w_1}}{\eta - \eta^{w_1}}[\eta Q + (1-\eta)S] + \frac{\eta - \eta^{w_2}}{\eta - \eta^{w_1}}\hat{S}.$$

Note that $\eta P+(1-\eta)R, \eta Q+(1-\eta)S, \hat{R}, \hat{S} \in [T^3, T^4]$, and $(\eta P+(1-\eta)R)_2 = \hat{R}_2, (\eta Q+(1-\eta)S)_2 = \hat{S}_2$. By assumption and Lemma A.2.9, the independence property holds for $(\eta P + (1-\eta)R, \eta Q + (1-\eta)S, \hat{R}, \hat{S})$. Whenever $\eta P + (1-\eta)R \not\sim \eta Q + (1-\eta)S$, I know $\hat{P} \not\sim \hat{Q}$, leading to a contradiction. Thus, $\eta P + (1-\eta)R \sim \eta Q + (1-\eta)S$.

The proof for the case with $\eta^{w_1} > \eta > 0$ is symmetric. Hence for all $\eta \in (0,1)$, $\eta P + (1-\eta)R \sim \eta Q + (1-\eta)S$. ■

The next Lemma A.2.11 shows (P, Q, R, S) with $P_2 = R_2$ and $Q_2 = S_2$ satisfies the independence property if the utilities of P and R are "sufficiently close".

Lemma A.2.11 *In Case 2, suppose $z = z_0$ or z_0'. Then a proper tuple $(P, Q, R, S) \in (\mathcal{L}^0(X_1) \times X_2)^4$ satisfies the independence property if $P \sim Q, R \sim S$, $P_2 = R_2 = y_1, Q_2 = S_2 = y_2$ with with $y_1, y_2 \in \Sigma^2(z)$ and there exists $x_1, x_2 \in X_1$ such that $P, Q, R, S \in \Gamma_{2,y_1}(x_1|z) \cap \Gamma_{2,y_2}(x_2|z)$.*

Proof of Lemma A.2.11. Suppose $(P, Q, R, S) \in (\mathcal{L}^0(X_1) \times X_2)^4$ is a proper tuple that satisfies the conditions stated in the lemma. By Lemma A.2.4, there exist $P_1', R_1' \in \Pi^1(x_1|z)$ and $Q_1', S_1' \in \Pi^1(x_2|z)$ such that $P' := (P_1', P_2) \sim P \sim Q \sim Q' := (Q_1', Q_2)$ and $R' := (R_1', R_2) \sim R \sim S \sim S' := (S_1', S_2)$. By (ii) of Lemma A.2.3, for any $\alpha \in (0,1)$, $\alpha P' + (1-\alpha)R' \sim \alpha Q' + (1-\alpha)S'$. Lemma A.2.1 implies $\alpha P'+(1-\alpha)R' \sim \alpha P+(1-\alpha)R$ and $\alpha Q'+(1-\alpha)S' \sim \alpha Q+(1-\alpha)S$. By transitivity of $\succsim$, $\alpha P + (1-\alpha)R \sim \alpha Q + (1-\alpha)S$. ■

By Lemma A.2.10, I extend the local property in Lemma A.2.11 to a global one.

Lemma A.2.12 *In Case 2, suppose $z = z_0$ or z_0'. Then a proper tuple $(P, Q, R, S) \in (\mathcal{L}^0(X_1) \times X_2)^4$ satisfies the independence property if $P_2 = R_2 = y_1, Q_2 = S_2 = y_2$ with $y_1, y_2 \in \Sigma^2(z)$.*

Proof of Lemma A.2.12. The independence property holds for (P, Q, R, S) trivially if $y_1 = y_2$. Without loss of generality, I assume $y_1 > y_2$ and $P \succ R$. Fix $\hat{T} \in \mathcal{L}^0(X_1) \times \{y_2\}$ and $a \in X_1$ with $\hat{T} \succ (a, y_1) \succ (a, y_2)$. Take any $T_1 = (z_1, y_2), T^2 = (z_2, y_2)$ with $\hat{T} \succ T^1 \succ T^2 \succ (a, y_1)$. As $y_1, y_2 \in \Sigma^2(z)$, by Lemma A.2.1, I can find $\hat{x}_1$, $\hat{x}_2 \in X_1^o$ and $p_1, q_1 \in \Pi^1(\hat{x}_1|z)$, $p_2, q_2 \in \Pi^1(\hat{x}_2|z)$ such that

$$(q_1, y_1) \sim T^1 \prec (p_1, y_1) \text{ and } (q_2, y_2) \sim T^1 \prec (p_2, y_2).$$

By Lemma A.2.4, I know

$$[T^1, (p_1, y_1)] \cap [T^1, (p_2, y_2)] \subseteq \Gamma(\hat{x}_1, y_1) \cap \Gamma(\hat{x}_2, y_2).$$

For any $z' \in [z_2, z_1]$, I can choose $\lambda_{z'}$ and $\eta_{z'} \in [0, 1]$ such that

$$\lambda_{z'}(q_1, y_1) + (1 - \lambda_{z'})(a, y_1) \sim (z', y_2) \sim \eta_{z'}(q_2, y_2) + (1 - \eta_{z'})(a, y_2).$$

By Lemma A.2.1, $\lambda_{z'}(p_1, y_1) + (1 - \lambda_{z'})(a, y_1) \succ \lambda_{z'}(q_1, y_1) + (1 - \lambda_{z'})(a, y_1)$ and $\eta_{z'}(p_2, y_2) + (1 - \eta_{z'})(a, y_2) \succ \eta_{z'}(q_2, y_2) + (1 - \eta_{z'})(a, y_2)$. Denote $\hat{x}_1^{z'}$, $\hat{x}_2^{z'}$ for each $z' \in [z_1, z_2]$ with

$$\lambda_{z'} q_1 + (1 - \lambda_{z'})\delta_a \in \Pi^1(\hat{x}_1^{z'}|z), \text{ and } \eta_{z'} q_2 + (1 - \eta_{z'})\delta_a \in \Pi^1(\hat{x}_2^{z'}|z).$$

This leads to

$$[(z', y_2), (\lambda_{z'} p_1 + (1 - \lambda_{z'})\delta_a, y_1)] \cap [(z', y_2), (\eta_{z'} p_2 + (1 - \eta_{z'})\delta_a, y_2)] \subseteq \Gamma_{2,y_1}(\hat{x}_1^{z'}|z) \cap \Gamma_{2,y_2}(\hat{x}_2^{z'}|z).$$

Take the union across all z' between z_1 and z_2, and by Continuity, I have

$$[T^2, T^1] \subseteq \bigcup_{z_2 \leq z' \leq z_1} \Big(\Gamma_{2,y_1}(\hat{x}_1^{z'}|z) \cap \Gamma_{2,y_2}(\hat{x}_2^{z'}|z) \Big). \tag{A.3}$$

In order to get an open cover of $[T^2, T^1]$, for $\varepsilon > 0$ small enough with $\hat{T} \succ (z_1 + \varepsilon, y_2) \succ (z_2 - \varepsilon, y_2) \succ (a, y_1)$, I have

$$[T^2, T^1] \subseteq \bigcup_{z_2 - \varepsilon \leq z' \leq z_1 + \varepsilon} \Big(\Gamma_{2,y_1}(\hat{x}_1^{z'}|z) \cap \Gamma_{2,y_2}(\hat{x}_2^{z'}|z) \Big).$$

For each $z_2 - \varepsilon \leq z' \leq z_1 + \varepsilon$, $\Gamma_{2,y_1}(\hat{x}_1^{z'}|z) \cap \Gamma_{2,y_2}(\hat{x}_2^{z'}|z)$ has a non-empty interior. Hence I can find an open cover of $[T^2, T^1] = [(z_2, y_2), (z_1, y_2)]$ as $\{C^{z'}\}_{z_2 - \varepsilon \leq z' \leq z_1 + \varepsilon}$ with $C^{z'} \subset \Gamma_{2,y_1}(\hat{x}_1^{z'}|z) \cap \Gamma_{2,y_2}(\hat{x}_2^{z'}|z)$. Notice that $X_1 \times \{y_2\}$ is isomorphic to $X_1 \subseteq \mathbb{R}$ and $[T^2, T^1] = [(z_2, y_2), (z_1, y_2)]$ is isomorphic to $[z_2, z_1]$, which is a closed and bounded interval. By Heine–Borel theorem, I can find a finite subcover of $\{C^{z'}\}_{z_2 - \varepsilon \leq z' \leq z_1 + \varepsilon}$ for $[T^2, T^1]$. Denote the subcover as $\{C^{z_k}\}_{k=1}^K$ and I can assume that $C^{z_k} \cap C^{z_{k+1}} \neq \emptyset$ for each $k = 1, ..., K-1$.

Take any proper tuple (P, Q, R, S) with $P \sim Q, R \sim S$, $P_2 = R_2 = y_1, Q_2 = S_2 = y_2$ with $y_1 > y_2 \in \Sigma^2(z)$. By Lemma A.2.11, the independence property holds for (P, Q, R, S) if $P, Q, R, S \in C^{z_k}$ for any $k = 1, ..., K$. Then Lemma A.2.10 implies that the independence property holds for (P, Q, R, S) if $P, Q, R, S \in [T^2, T^1] \subseteq \bigcup_{k=1}^K C^{z_k}$.

Recall that I can denote $X_1 = [\underline{c}_1, \overline{c}_1]$ where $\underline{c}_1, \overline{c}_1 \in \mathbb{R}$. By arbitrariness of $\hat{T}, T^1$, T^2 and $a \in X_2$, the given tuple (P, Q, R, S) satisfies the independence property if $(\overline{c}_1, y_2) \succ P \sim Q \succ R \sim S \succ (\underline{c}_1, y_1)$.

Now suppose $(\overline{c}_1, y_2) \succ P$ and $R \sim (\underline{c}_1, y_1)$. By Monotonicity, $R = (\underline{c}_1, y_1)$. Take a sequence of $\{\lambda_n\}_{n \geq 1} \subset (0, 1)$ with $\lambda_n \to 0$. For each n, denote $S^n = (\lambda_n Q_1 + (1 - \lambda_n) S_1, S_2)$ and by Lemma A.2.1, I can find β_n with $R^n = (\beta_n P_1 + (1 - \beta_n) R_1, R_2) \sim S^n$. Clearly, $R^n \succ R = (\underline{c}_1, y_1)$ for each n and hence the independence property holds for (P, Q, R^n, S^n), that is, for each $\alpha \in (0, 1)$,

$$\alpha P + (1 - \alpha) R^n \sim \alpha Q + (1 - \alpha) S^n.$$

As n goes to infinity, β_n converges to 0 and hence $S^n \xrightarrow{w} S, R^n \xrightarrow{w} R$. By continuity of $\succsim$ on $\hat{\mathcal{P}}$, I have for each $\alpha \in (0, 1)$,

$$\alpha P + (1 - \alpha) R \sim \alpha Q + (1 - \alpha) S.$$

A similar proof works for the case where $P \sim (\overline{c}_1, y_2)$ and/or $R \sim (\underline{c}_1, y_1)$. Hence, the independence property holds for all (P, Q, R, S) with $P \sim Q, R \sim S$, $P_2 = R_2 = y_1, Q_2 = S_2 = y_2$ with $y_1, y_2 \in \Sigma^2(z)$ for $z \in \{z_0, z_0'\}$. ■

Lemma A.2.13 *In Case 2, a proper tuple $(P, Q, R, S) \in (\mathcal{L}^0(X_1) \times X_2)^4$ satisfies the independence property if $P_2 = R_2 = y_1, Q_2 = S_2 = y_2$ with $y_1, y_2 \in X_2$.*

Proof of Lemma A.2.13. By Lemma A.2.7, $\Sigma^2(z_0) \cup \Sigma^2(z_0') = X_2$ and $\Sigma^2(z_0) \cap \Sigma^2(z_0') \neq \emptyset$. By Lemma A.2.12, it suffices to consider the case where $y_1 \in \Sigma^2(z_0) \backslash \Sigma^2(z_0')$ and $y_2 \in \Sigma^2(z_0') \backslash \Sigma^2(z_0)$. Without loss of generality, assume $y_1 > y_2$ and $P \sim Q \succ R \sim S$.

Since $\Sigma^2(z_0)$ and $\Sigma^2(z_0')$ are open, there exists $z \in \Sigma^2(z_0) \cap \Sigma^2(z_0')$ with $y_1 > z > y_2$. By Monotonicity and Lemma A.2.1, I can find $P' = (P_1', z)$ and $R' = (R_1', z)$ such that $P' \sim P \sim Q, R' \sim R \sim S$. By Lemma A.2.12, the independence property holds for (P, P', R, R') and (Q, P', S, R') respectively. Hence for any $\alpha \in (0, 1)$,

$$\alpha P + (1-\alpha)R \sim \alpha P' + (1-\alpha)R' \sim \alpha Q + (1-\alpha)S.$$

■

I am now ready to show that $\succsim$ can be represented by V^{ANB-CN}.

Recall that I can focus on the restriction of $\succsim$ to $\mathcal{L}^0(X_1) \times X_2$. For any (p_1, x), $(p_2, x) \in \mathcal{L}^0(X_1) \times X_2$ with $(p_1, x) \succ (p_2, x)$, I claim that there exists some function f representing $\succsim$ on $(\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$ such that f is continuous and linear in the first dimension, that is, for any $(q_1, y), (q_2, y) \in (\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$ and $\alpha \in [0, 1]$, $f(\alpha q_1 + (1-\alpha)q_1, y) = \alpha f(q_1, y) + (1-\alpha) f(q_2, y)$. In addition, f is unique up to a positive affine transformation. I call such f a MAP_1 function.

To prove the claim, notice that by Lemma A.2.3, for any $Q \in [(p_2, x), (p_1, x)]$, there exists a unique $\lambda_Q \in [0, 1]$ such that $Q \sim \lambda_Q (p_1, x) + (1 - \lambda_Q)(p_2, x)$. Define $f : [(p_2, x), (p_1, x)] \to [0, 1]$ such that $f(Q) = \lambda_Q$. As $(p_1, x) \succ (p_2, x)$, I have

$$f(Q) \geq f(Q') \Longleftrightarrow Q \sim \lambda_Q (p_1, x) + (1-\lambda_Q)(p_2, x) \succsim \lambda_{Q'} (p_1, x) + (1-\lambda_{Q'})(p_2, x) \sim Q'.$$

Hence f represents $\succsim$ on $[(p_2, x), (p_1, x)]$. Continuity of $\succsim$ on $\hat{\mathcal{P}}$ assures that f is continuous. Then I show the linearity of f in the first dimension on $(\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$. Take any $(q_1, y), (q_2, y) \in (\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$. By definition of f and Lemma A.2.1, I have

$$\begin{aligned}(q_1, y) &\sim f(q_1, y)(p_1, x) + (1 - f(q_1, y))(p_2, x),\\ (q_2, y) &\sim f(q_2, y)(p_1, x) + (1 - f(q_2, y))(p_2, x).\end{aligned}$$

Clearly, for any $\alpha \in (0, 1)$, $\alpha(q_1, y) + (1-\alpha)(q_2, y) \in (\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$. By definition of f and Lemma A.2.12,

$$\begin{aligned}\alpha(q_1, y) + (1 - \alpha)(q_2, y) \sim &[\alpha f(q_1, y) + (1 - \alpha) f(q_2, y)](p_1, x)\\ &+ [1 - \alpha f(q_1, y) - (1 - \alpha) f(q_2, y)](p_2, x)\\ \sim &f(\alpha q_1 + (1 - \alpha) q_2, y)(p_1, x)\\ &+ (1 - f(\alpha q_1 + (1 - \alpha) q_2, y))(p_2, x)\end{aligned}$$

By Lemma A.2.1 given x in source 2, I know $f(\alpha q_1 + (1-\alpha)q_2, y) = \alpha f(q_1, y) + (1 - \alpha) f(q_2, y)$. Easy to see that a positive affine transformation of f also represents $\succsim$ on $(\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$.

To see that f is unique up to a positive affine transformation, assume f, g represent $\succsim$ on $(\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$ and they are continuous and linear in the first dimension. Without loss of generality, let $f(p_2, x) = g(p_2, x), f(p_1, x) = g(p_1, x)$. Recall that for any $Q \in (\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$, there is a unique λ_Q with $Q \sim \lambda_Q(p_1, x) + (1 - \lambda_Q)(p_2, x)$. By linearity of f and g in the first dimension on $(\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$, I have

$$\begin{aligned}f(Q) =&\lambda_Q f(p_1, x) + (1 - \lambda_Q) f(p_2, x)\\ =&\lambda_Q f(p_1, x) + (1 - \lambda_Q) f(p_2, x)\\ =&g(Q)\end{aligned}$$

Hence $f \equiv g$ on $(\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$.

Now suppose that $\succsim$ admits a MAP_1 representation on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^2, T^1]$ and $(\mathcal{L}^0(X_1) \times X_2) \cap [T^4, T^3]$ where $T_2^1 = T_2^2 = x, T_2^3 = T_2^4 = y \neq x$. If $[T^2, T^1] \cap [T^4, T^3]$ has an empty interior, then it is trivial to show that $\succsim$ admits a MAP_1 representation on $(\mathcal{L}^0(X_1) \times X_2) \cap \Big([T^2, T^1] \cup [T^4, T^3]\Big)$. Now assume without loss of generality that $T^1 \succ T^3 \succ T^2 \succ T^4$. Denote the MAP_1 function on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^2, T^1]$ as f_x and the MAP_1 function on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^4, T^3]$ as f_y. Then f_x must be a positive affine transformation of f_y on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^2, T^3]$. Fix f_x and let $f_y(T^i) = f_x(T^i)$, $i = 2, 3$, then I have $f_x = f_y$ on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^2, T^3]$. Define $\hat{f} = f_x$ on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^2, T^1]$ and $\hat{f} = f_y$ on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^4, T^3]$. Then easy to show that $\hat{f}$ is a MAP_1 function on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^4, T^1]$ and is unique up to a positive affine transformation.

By a similar argument, if $\succsim$ admits a MAP_1 representation on $(\mathcal{L}^0(X_1) \times X_2) \cap [T^2, T^1]$ and $(\mathcal{L}^0(X_1) \times X_2) \cap [T^4, T^3]$, then it also admits a MAP_1 representation on $(\mathcal{L}^0(X_1) \times X_2) \cap \Big([T^2, T^1] \cup [T^4, T^3]\Big)$ and the utility function is unique up to a positive affine transformation. Recall that $\succsim$ admits a MAP_1 representation on $(\mathcal{L}^0(X_1) \times X_2) \cap [(p_2, x), (p_1, x)]$. By varying $p_1, p_2 \in \mathcal{L}^0(X_1)$ and $x \in X_2$, I can show that there exists a MAP_1 function f^* representing $\succsim$ on $\mathcal{L}^0(X_1) \times X_2$ and it is unique up to a positive affine transformation.

Define $w(x, y) = f^*(x, y)$ for all $(x, y) \in X_1 \times X_2$. By construction of f^*, w is continuous. By Monotonicity, w is strictly monotonic. Recall that $\succsim$ satisfies Correlation Neutrality and for any $(p, q) \in \hat{\mathcal{P}}$, $(p, q) \succsim (p, CE_{v_2}(q))$ for some continuous and strictly monotonic v_2. Hence for any $P, Q \in \mathcal{P}$,

$$\begin{aligned} P \succsim Q &\iff (P_1, CE_{v_2}(P_2)) \succsim (Q_1, CE_{v_2}(Q_2)) \\ &\iff \sum_x w(x, CE_{v_2}(P_2)) P_1(x) \geq \sum_y w(y, CE_{v_2}(Q_2)) Q_1(y) \\ &\iff V^{ANB-CN}(P) \geq V^{ANB-CN}(Q). \end{aligned}$$

Thus, in Case 2, $\succsim$ is represented by V^{ANB-CN}.

Case 3: There exists $(z_1, z_2) \in X_1 \times X_2$ **such that** $(p, q) \sim (p', q)$ **for all** $(x, y) \in X_1 \times X_2$**,** $p, p' \in \Pi^1(x|z_2)$ **and** $q \in \Pi^2(y|z_1)$**, and** $\exists (x_0, y_0) \in X_1 \times X_2$**,** $p_0 \in \Pi^1(x_0|z_2)$ **and** $q_0, q_0' \in \Pi^2(y_0|z_1)$ **with** $(p_0, q_0) \succ (p_0, q_0')$**.**

This is symmetric to Case 2 and $\succsim$ can be represented by

$$\hat{V}(P) = \sum_y w(CE_{v_1}(P_1), y)P_2(y)$$

where v_1 and w are continuous and strictly monotonic. However, I will show that given part (iii) of Independence*, $\succsim$ actually admits a CN representation.

Since $\succsim$ is represented by $\hat{V}$, I can focus on the restriction of $\succsim$ on $X_1 \times \mathcal{L}^0(X_2)$ and a proper tuple $(P, Q, R, S) \in \hat{\mathcal{P}}^4$ satisfies the independence property if $P_1 = R_1, Q_1 = S_1$.

By Lemma A.2.9, it suffices to consider the tuple with $P \sim Q, R \sim S$. Fix $y_0 \in X_2^o$ and let $P_2 = R_2 = y_0$. Denote $\hat{v}_{1|y_0}(x) = w(x, y_0)$ for each $x \in X_1$. I want to show that $\hat{v}_{1|y_0}$ is a positive affine transformation of v_1. Define $\Sigma^1(z) := \{x \in X_1 : \exists\ y \in X_2, q, q' \in \Pi^2(y|z) \text{ s.t. } (x, q) \succ (x, q')\}$. By Lemma A.2.7, there exists $z_0, z_0' \in X_1$ such that $\Sigma^1(z_0) \cup \Sigma^1(z_0') = X_1$ and $\Sigma^1(z_0) \cap \Sigma^1(z_0') \neq \emptyset$.

Now consider $Q \sim P$ such that $Q_1 = x_0 \in X_1^0$. Then $x_0 \in \Sigma^1(z)$ for some $z \in \{z_0, z_0'\}$ and there exists $h(x_0) \in X_1$ with $x_0 < h(x_0)$ and $[(x_0, y_0), (h(x_0), y_0)] \subset \Gamma_{1,x_0}(y_0|z)$. By Lemma A.2.4 and Lemma A.2.3, for any $\alpha \in (0, 1)$ and $P \sim Q \in [(x_0, y_0), (h(x_0), y_0)]$, let $R = S = (x_0, y_0)$, then $\alpha P + (1 - \alpha)R = \alpha Q + (1 - \alpha)R$. That is, for any $x \in [x_0, h(x_0)]$,

$$\alpha \hat{v}_{1|y_0}(x) + (1 - \alpha)\hat{v}_{1|y_0}(x_0) = \hat{v}_{1|y_0}(v_1^{-1}(\alpha v_1(x) + (1 - \alpha)v_1(x_0))).$$

Normalize that $\hat{v}_{1|y_0}(x_0) = v_1(x_0)$ and $\hat{v}_{1|y_0}(h(x_0)) = v_1(h(x_0))$. By continuity and monotonicity of v_1 and $\hat{v}_{1|y_0}$, $\hat{v}_{1|y_0}(x) = v_1(x)$ for all $x \in [x_0, h(x_0)]$.

When x_0' is close enough to x_0, $[x_0, h(x_0)]] \cap [x_0', h(x_0')]$ has nonempty interior and it is easy to show $\hat{v}_{1|y_0}(x) = v_1(x)$ for all $x \in [x_0, h(x_0)] \cup [x_0', h(x_0')]$. Then I vary x_0 across X_1^0 and conclude that $\hat{v}_{1|y_0}(x) = v_1(x)$ for all $x \in X_1^0$. By continuity of $\hat{v}_{1|y_0}$

and v_1, $\hat{v}_{1|y_0}(x) = v_1(x)$ for all $x \in X_1$. Since $y_0 \in X_2^o$ is arbitrarily selected and by continuity of w, for each $y \in X_2$ $\hat{v}_{1|y} = w(\cdot, y)$ is a positive affine transformation of v_1. Hence for each $P \in \mathcal{P}$, $\hat{V}(P)$ can be rewritten as

$$\hat{V}(P) = \sum_y w(CE_{v_1}(P_1), y)P_2(y) = \sum_{x,y} w(x,y)P_1(x)P_2(y).$$

Thus, $\succsim$ admits a CN representation.

Case 4: For any $(z_1, z_2) \in X_1 \times X_2$, there exist $x_1, x_2 \in X_1, y_1, y_2 \in X_2$ and $(p_1, q_1), (p_1, q_1') \in \Pi^1(x_1|z_2) \times \Pi^2(y_1|z_1)$, $(p_2, q_2), (p_2', q_2) \in \Pi^1(x_2|z_2) \times \Pi^2(y_2|z_1)$ such that $(p_1, q_1) \succ (p_1, q_1')$ and $(p_2, q_2) \succ (p_2', q_2)$.

For each $(z_1, z_2) \in X_1 \times X_2$, denote

$$\hat{\Sigma}^1(z_1, z_2) := \{x \in X_1 : \exists\ y \in X_2 \text{ and } p \in \Pi^1(x|z_2), q, q' \in \Pi^2(y|z_1) \text{ s.t. } (p,q) \succ (p,q')\},$$
$$\hat{\Sigma}^2(z_1, z_2) := \{y \in X_2 : \exists\ x \in X_1 \text{ and } p, p' \in \Pi^1(x|z_2), q \in \Pi^2(y|z_1) \text{ s.t. } (p,q) \succ (p',q)\}.$$

Easy to see that $\hat{\Sigma}^1(z_1, z_2)$ and $\hat{\Sigma}^1(z_1, z_2)$ are open and $\hat{\Sigma}^1(z_1, z_2) \neq \emptyset$ and $\hat{\Sigma}^2(z_1, z_2) \neq \emptyset$. Furthermore, the following lemma holds.

Lemma A.2.14 *In Case 4, suppose $(z_1, z_2) \in X_1 \times X_2$. Then for any $(x, y) \in (\hat{\Sigma}^1(z_1, z_2) \cap X_1^o) \times (\hat{\Sigma}^2(z_1, z_2) \cap X_2^o)$, there exist $P, Q, R, S \in \Pi^1(x|z_2) \times \Pi^2(y|z_1)$ such that $P_1 = Q_1$, $R_2 = S_2$, $P \succ Q$ and $R \succ S$.*

Proof of Lemma A.2.14. By definition, there exist $x' \in X_1, y' \in X_2$ and $(p_1, q_1), (p_1, q_1') \in \Pi^1(x|z_2) \times \Pi^2(y'|z_1)$, $(p_2, q_2), (p_2', q_2) \in \Pi^1(x'|z_2) \times \Pi^2(y|z_1)$ such that $(p_1, q_1) \succ (p_1, q_1')$ and $(p_2, q_2) \succ (p_2', q_2)$. By Lemma A.2.1, given p_1 in source 1, as $y \in X_2^o$, I can find $\hat{q}_1, \hat{q}_1' \in \Pi^2(y|z_1)$ with $(p_1, \hat{q}_1) \succ (p_1, \hat{q}_1')$. Similarly, given q_2 in source 2, as $x \in X_1^o$, I can find $\hat{p}_2, \hat{p}_2' \in \Pi^1(x_1|z_2)$ with $(\hat{p}_2, q_2) \succ (\hat{p}_2', q_2)$. This completes the proof. ■

For any $x_1 \in \hat{\Sigma}^1(z_1, z_2)$ and $x_2 \in \hat{\Sigma}^2(z_1, z_2)$, I denote

$$\hat{\Pi}^1(x_1|z_1, z_2) = \left\{p \in \Pi^1(x_1|z_2) : \exists\ x_2' \in X_2, q \in \Pi^2(x_2'|z_1) \text{ s.t. } (p,q) \not\sim (p, x_2')\right\},$$

$$\hat{\Pi}^2(x_2|z_1,z_2)=\Big\{q\in\Pi^2(x_2|z_1):\exists\ x_1'\in X_1, p\in\Pi^1(x_1'|z_2)\text{ s.t. }(p,q)\not\succ(x_1',q)\Big\}.$$

By definition $\hat{\Pi}^i(x_i|z_1,z_2)\subseteq\Pi^i(x_i|z_{-i})$ for $i=1,2$.

Lemma A.2.15 *For each $i=1,2$ and $x_i\in\Sigma^i$, $cl(\hat{\Pi}^i(x_i|z_1,z_2))=\Pi^i(x_i|z_{-i})$.*

Proof of Lemma A.2.15. I will prove the result for $i=1$. The proof for $i=2$ is symmetric. By definition of $x_1\in\hat{\Sigma}^1(z_1,z_2)$, I can find $p\in\hat{\Pi}^1(x_1|z_1,z_2)$, $x_2'\in X_2$ and $q,q'\in\Pi^2(x_2'|z_1)$ such that $(p,q)\succ(p,q')$. For any $p^o\in\Pi^1(x_1|z_2)\backslash\hat{\Pi}^1(x_1|z_1,z_2)$, I have $(p^o,q)\sim(p^o,q')$. By Lemma A.2.3, the independence property holds for $((p,q),(p,q'),(p^o,q),(p^o,q'))$ as $p,p^o\in\Pi^1(x_1|z_2)$. Then for any $\lambda\in(0,1)$, $(\lambda p+(1-\lambda)p^o,q)\succ(\lambda p+(1-\lambda)p^o,q')$, which implies $\lambda p+(1-\lambda)p^o\in\hat{\Pi}^1(x_1|z_1,z_2)$. Let $\lambda\to0$ and I have $p^o\in cl(\hat{\Pi}^1(x_1|z_1,z_2))$. ■

Now I introduce the local independence property similar to Lemma A.2.11. By Lemma A.2.9, it suffices to focus on the tuple (P,Q,R,S) with $P\sim Q$ and $R\sim S$.

Lemma A.2.16 *In Case 4, suppose $(z_1,z_2)\in X_1\times X_2$. Then a proper tuple (P,Q,R,S) satisfies the independence property if there exist some $i,j\in\{1,2\}$, $x_i\in\hat{\Sigma}^i(z_1,z_2)\cap X_i^o$, $x_j\in\hat{\Sigma}^j(z_1,z_2)\cap X_j^o$ and $y\in X_{-i}$, $y'\in X_{-j}$ such that $P_i=R_i=p\in\hat{\Pi}^i(x_i|z_1,z_2)$, $Q_j=S_j=q\in\hat{\Pi}^j(x_j|z_1,z_2)$, $P\sim Q$, $R\sim S$ and $P,R\in\Gamma_{i,p}(y|z_i)\cap\Gamma_{j,q}(y'|z_j)$.*

Proof of Lemma A.2.16. By Lemma A.2.4, one can find P',R' with $P_i'=R_i'=p$ and $P_{-i}',R_{-i}'\in\Pi^{-i}(y|z_i)$ such that $P\sim P'$ and $R\sim R'$. By Lemma A.2.1, for any $\lambda\in(0,1)$, $\lambda P+(1-\lambda)R\sim\lambda P'+(1-\lambda)R'$. Similarly, I can find Q',S' with $Q_j'=S_j'=q$, $Q_{-j}',R_{-j}'\in\Pi^{-j}(y'|z_j)$ and $\lambda Q+(1-\lambda)S\sim\lambda Q'+(1-\lambda)S'$ for any $\lambda\in(0,1)$. Easy to check that the conditions in (ii) of Lemma A.2.3 hold for the tuple (P',Q',R',S') and hence for any $\lambda\in(0,1)$, $\lambda P+(1-\lambda)R\sim\lambda P'+(1-\lambda)R'\sim\lambda Q'+(1-\lambda)S'\sim\lambda Q+(1-\lambda)S$. ■

The independence property in Lemma A.2.16 is non-trivial, since by definition of $\hat{\Sigma}^i(z_1,z_2)$ and $\hat{\Pi}^i(x_i|z_1,z_2)$ and Lemma A.2.14, for each $i,j\in\{1,2\}$, there exists $P,Q,R,S\in\Gamma_{i,p}(y|z_i)\cap\Gamma_{j,q}(y'|z_j)$ such that $P\sim Q\succ R\sim S$ and $P_i=R_i=p,Q_j=$

$S_j = q$. Like in Case 2, my goal is to extend the local property in Lemma A.2.16. The arguments are essentially the same as Lemma A.2.10, Lemma A.2.12 and Lemma A.2.13 in Case 2. Here I just summarize the main steps.

- Step 1. Fix $(z_1, z_2) \in X_1 \times X_2$. Consider a proper tuple (P, Q, R, S) where there exist some $i, j \in \{1, 2\}$, $x_i \in \hat{\Sigma}^i(z_1, z_2) \cap X_i^o$ and $x_j \in \hat{\Sigma}^j(z_1, z_2) \cap X_j^o$ such that $P_i = R_i = p \in \hat{\Pi}^i(x_i|z_1, z_2)$, $Q_j = S_j = q \in \hat{\Pi}^j(x_j|z_1, z_2)$, $P \sim Q, R \sim S$,
  - Step 1.1. Use arguments in Lemma A.2.10, Lemma A.2.12 to show that the independence property holds if $P, R \in \Gamma_{j,q}(y'|z_j)$ for some $y' \in X_{-j}$.
  - Step 1.2. Again use arguments in Lemma A.2.10, Lemma A.2.12 to show that the independence property always holds.
- Step 2. Fix $(z_1, z_2) \in X_1 \times X_2$. By Continuity and Lemma A.2.15, a proper tuple (P, Q, R, S) satisfies the independence property if there exist some $i, j \in \{1, 2\}$, $x_i \in \hat{\Sigma}^i(z_1, z_2)$ and $x_j \in \hat{\Sigma}^j(z_1, z_2)$ such that $P_i = R_i = p \in \Pi^i(x_i|z_{-i})$, $Q_j = S_j = q \in \Pi^j(x_j|z_{-j})$ and $P \sim Q, R \sim S$.
- Step 3. Now vary $(z_1, z_2) \in X_1 \times X_2$. By a similar argument in Lemma A.2.7, one can find $(z_1', z_2'), (z_1'', z_2'') \in X_1 \times X_2$ such that

$$\begin{aligned}
&\hat{\Sigma}^1(z_1', z_2') \cup \hat{\Sigma}^1(z_1'', z_2') = \hat{\Sigma}^1(z_1', z_2'') \cup \hat{\Sigma}^1(z_1'', z_2'') = X_1,\\
&\hat{\Sigma}^1(z_1', z_2') \cup \hat{\Sigma}^1(z_1', z_2'') = \hat{\Sigma}^1(z_1'', z_2') \cup \hat{\Sigma}^1(z_1'', z_2'') = X_2,\\
&\hat{\Sigma}^1(z_1', z_2') \cap \hat{\Sigma}^1(z_1'', z_2') \neq \emptyset \ , \ \hat{\Sigma}^1(z_1', z_2'') \cap \hat{\Sigma}^1(z_1'', z_2'') \neq \emptyset,\\
&\hat{\Sigma}^1(z_1', z_2') \cap \hat{\Sigma}^1(z_1', z_2'') \neq \emptyset \ , \ \hat{\Sigma}^1(z_1'', z_2') \cap \hat{\Sigma}^1(z_1'', z_2'') \neq \emptyset.
\end{aligned}$$

  Then by the same arguments in Lemma A.2.13, I can show that a proper tuple (P, Q, R, S) satisfies the independence property if there exist some $i, j \in \{1, 2\}$ such that $P_i = R_i$, $Q_j = S_j$ and $P \sim Q, R \sim S$.

Now I will show that $\succsim$ can be represented by some V^{CN}. By Lemma A.2.9 and Step 3 above, for any $P, Q, R, S \in \hat{\mathcal{P}}$, if $P_i = R_i$, $Q_j = S_j$ for some $i, j \in \{1, 2\}$, $R \sim S$ and $\alpha \in (0, 1)$, then $P \succsim Q \iff \alpha P + (1-\alpha)R \succsim \alpha Q + (1-\alpha)S$.

Apply Theorem 1 in Chapter 7.2 of Fishburn (1982) to $\hat{\mathcal{P}}$, and there exists a continuous and multilinear[1] function V representing $\succsim$ on $\hat{\mathcal{P}}$, which is unique up to a positive affine transformation. By Correlation Neutrality, define V' on $\mathcal{P}$ such that $V'(P) = V(P_1, P_2)$ for each $P \in \mathcal{P}$. Then V' represents $\succsim$ on $\mathcal{P}$.

I claim that V' is a CN representation. To see this, define $w(x_1, x_2) = V(x_1, x_2)$. Since V is continuous, w is continuous. Also by Monotonicity, w is continuous and strictly monotonic. Moreover, since V is multilinear and each lottery has a finite support, for each $P \in \mathcal{P}$,

$$V'(P) = V(P_1, P_2) = \sum_{x,y} V(x, y) P_1(x) P_2(y) = \sum_{x,y} w(x, y) P_1(x) P_2(y).$$

Thus in Case 4, $\succsim$ is represented by some $V^{CN} = V'$.

To summarize, if $\succsim$ satisfies Weak Order, Monotonicity, Continuity, Independence* and Correlation Neutrality, then $\succsim$ an be represented by one of the following types of utility functions: V^{CN} (Case 3 and Case 4), V^{ANB-CN} (Case 2) or V^{FNB} (Case 1).

A.3 Proof of Theorem 1.4.2

Necessity. Since Correlation Consistency is weaker than Correlation Neutrality, by Theorem 1.4.1, V^{CN}, V^{ANB-CN} and V^{FNB} satisfy all the axioms. Since w is continuous and strictly monotonic, V^{EU} also satisfy those axioms. Hence I can focus on V^{ANB} with

$$V^{ANB}(P) = \sum_{x} w(x, CE_v(P_{2|x})) P_1(x), \text{ for all } P \in \mathcal{P},$$

where w and v are continuous and strictly monotonic. This guarantees that V^{ANB} satisfies Weak Order and Monotonicity. Because V^{ANB} agrees with V^{ANB-CN} on $\hat{\mathcal{P}}$, V^{ANB} satisfies Continuity 3 and Independence*.

[1]Suppose $\mathcal{M}_1, \mathcal{M}_2 \subseteq \mathcal{L}^0(\mathbb{R})$ are mixture sets. A function V is multilinear on $\mathcal{M}_1 \times \mathcal{M}_2$ if $V(\lambda p + (1-\lambda)r, q) = \lambda V(p, q) + (1-\lambda)V(r, q)$ and $V(p, \lambda q + (1-\lambda)s) = \lambda V(p, q) + (1-\lambda)V(p, s)$ for all $\lambda \in (0, 1)$, $p, r \in \mathcal{M}_1$ and $q, s \in \mathcal{M}_2$.

To verify Continuity 1, for any $P, Q \in \mathcal{P}$ and $\lambda \in [0,1]$,

$$\begin{aligned} V^{ANB}(\lambda P + (1-\lambda)Q) =& \lambda \sum_{\substack{x: P_1(x)>0, \\ Q_1(x)=0}} w(x, CE_v(P_{2|x}))P_1(x) \\ &+ (1-\lambda) \sum_{\substack{x: Q_1(x)>0, \\ P_1(x)=0}} w(x, CE_v(Q_{2|x}))Q_1(x) \\ &+ \sum_{\substack{x: Q_1(x)>0, \\ P_1(x)>0}} w(x, CE_v(\alpha P_{2|x} + (1-\alpha)Q_{2|x}))[\lambda P_1(x) + (1-\lambda)Q_1(x)], \end{aligned}$$

where

$$\alpha = \frac{\lambda P_1(x)}{\lambda P_1(x) + (1-\lambda)Q_1(x)}.$$

Then $V^{ANB}(\lambda P + (1-\lambda)Q)$ is continuous in λ and $\succsim$ satisfies Continuity 1.

To verify Continuity 2, for ε_n small enough,

$$V^{ANB}(P_{\varepsilon_n}) = \sum_{x \in X_1} w\Big(\phi_{\varepsilon_n}(x,y)_1, v^{-1}\Big(\sum_{y \in X_2} v(\phi_{\varepsilon_n}(x,y)_2) P_{2|x}(y)\Big)\Big) P_1(x).$$

By definition, ϕ_ε is continuous in ε. By continuity of w and v_2, $V^{ANB}(P_{\varepsilon_n})$ converges to $V^{ANB}(P)$ as ε_n converges to $(0,0)$.

Finally, I verify Correlation Consistency. For each $P, Q, R, S \in \mathcal{P}$ and $\alpha \in (0,1)$, if $P_i = Q_i$ for $i = 1, 2$ and $\text{supp}(P_1) \cap \text{supp}(R_1) = \text{supp}(P_1) \cap \text{supp}(S_1) = \emptyset$, then

$$V^{ANB}(\alpha P + (1-\alpha)R) = \alpha V^{ANB}(P) + (1-\alpha)V^{ANB}(R),$$

$$V^{ANB}(\alpha Q + (1-\alpha)S) = \alpha V^{ANB}(Q) + (1-\alpha)V^{ANB}(S).$$

Hence $P \succ Q, R \sim S$ implies $\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$.

Sufficiency. First, I utilize Theorem 1.4.1 to derive implications of axioms on $\hat{\mathcal{P}}$. Formally, define $\hat{\succsim}$ on $\mathcal{P}$ which satisfies Correlation Neutrality and agrees with $\succsim$ on $\hat{\mathcal{P}}$. Clearly, $\hat{\succsim}$ satisfies Weak Order, Monotonicity, Continuity, Independence* and Correlation Consistency. Then Theorem 1.4.1 leads to the following corollary.

Corollary A.3.1 *$\hat{\succsim}$ can be represented by one of the following types of utility functions: V^{CN}, V^{ANB-CN} or V^{FNB}. Moreover, if $\succsim$ satisfies Correlation Neutrality, then $\succsim = \hat{\succsim}$.*

For the rest of the proof, I maintain the assumption that $\succsim$ violates Correlation Neutrality, that is, there exists $P \succ \tilde{P}$ with $P_i = \tilde{P}_i$ for $i = 1, 2$.

For any $(p, q) \in \hat{\mathcal{P}}$, denote $M(p, q)$ as the set of lotteries whose marginal lotteries are p and q respectively. For any $P, R \in \mathcal{P}$, I say P and R are *compatible*, or P is *compatible* with R if $\text{supp}(P_1) \cap \text{supp}(R_1) = \emptyset$. Easy to see that if R is compatible with both P and Q, then R is also compatible with $\lambda P + (1-\lambda)Q$ for any $\lambda \in (0, 1)$. If P is compatible with Q, then P is compatible with any $Q' \in M(Q_1, Q_2)$.

The first main challenge is that betweenness is not directly implied by the axioms, that is, for $\lambda \in (0, 1)$, it is not guaranteed that $P \succ \lambda P + (1-\lambda)\tilde{P} \succ \tilde{P}$. However, I have the following weaker and local version of the betweenness property.

Lemma A.3.2 *For any $Q \succ Q'$, there exists $Q^* = \lambda^* Q + (1-\lambda^*)Q'$ for some $\lambda^* \in [0, 1]$ such that for any $\varepsilon > 0$, there exists $\lambda_\varepsilon \in (\lambda^* - \varepsilon, \lambda^* + \varepsilon) \cap [0, 1]$ with $Q^* \not\sim \lambda_\varepsilon Q + (1-\lambda_\varepsilon)Q'$.*

Proof of Lemma A.3.2. Suppose the result fails. Then for any $\lambda \in [0, 1]$, there exists $\varepsilon_\lambda > 0$ such that for any $\lambda' \in (\lambda - \varepsilon_\lambda, \lambda + \varepsilon_\lambda) \cap [0, 1]$, $\lambda Q + (1-\lambda)Q' \sim \lambda' Q + (1-\lambda')Q'$. Notice that $\{(\lambda - \varepsilon_\lambda, \lambda + \varepsilon_\lambda)\}_{\lambda \in [0,1]}$ forms an open cover of the compact set $[0, 1]$. I can find a finite subcover of $[0, 1]$. By transitivity of $\succsim$, I know that $\lambda Q + (1-\lambda)Q' \sim \lambda' Q + (1-\lambda')Q'$ for all $\lambda, \lambda' \in [0, 1]$, which leads to $Q \sim Q'$, a contradiction. ■

The next lemma shows that if $R \sim S$ and Q is compatible with R, S, then the independence property holds.

Lemma A.3.3 *For any $\alpha \in (0, 1)$ and $Q, R, S \in \mathcal{P}$ with $R \sim S$ and Q compatible with R, S,*

$$\alpha Q + (1-\alpha)R \sim \alpha Q + (1-\alpha)S.$$

Proof of Lemma A.3.3. For $P \succ \tilde{P}$ with $\tilde{P} \in M(P_1, P_2)$. Recall that $X_1 = [\underline{c}_1, \bar{c}_1]$. If $P_1(\underline{c}_1) > 0$, then by Continuity 2, I can find $\varepsilon = (\varepsilon^1, 0)$ with $\varepsilon^1 > 0$ small enough such that $P_\varepsilon \succ \tilde{P}_\varepsilon$. Denote $P' = P_\varepsilon$ and $\tilde{P}' = \tilde{P}_\varepsilon$. Then $P' \succ \tilde{P}'$ with $\tilde{P}' \in M(P_1', P_2')$ and $P_1'(\underline{c}_1) = 0$. If $P_1'(\underline{c}_1) = 0$, repeat the above argument by choosing some $\varepsilon^1 < 0$ and I can construct $P'' \succ \tilde{P}''$ with $\tilde{P}'' \in M(P_1'', P_2'')$ and $P_1''(\underline{c}_1) = P_1''(\bar{c}_1) = 0$. Thus, without loss of generality, I assume $P_1(\underline{c}_1) = P_1(\bar{c}_1) = 0$, that is, $\text{supp}(P_1) \in X_1^o$.

Denote $P^* = \lambda^* P + (1-\lambda^*)\tilde{P}$ as the lottery found in Lemma A.3.2. Clearly, either $P^* \not\sim P$ or $P^* \not\sim \tilde{P}$. P_1, P_2 are not degenerate. By Lemma A.3.2, for any $n > 0$, there exists $\lambda_n \in (\lambda^* - 1/n, \lambda^* + 1/n) \cap [0,1]$ with $P^* \not\sim \lambda_n P + (1-\lambda_n)\tilde{P} := P^n$. By completeness, for each n, either $P^n \succ P^*$ or $P^* \succ P^n$. Then I can find a subsequence of $\{P^n\}_{n\geq 1}$, still denoted as $\{P^n\}_{n\geq 1}$, such that either $P^n \succ P^*$ for all n or $P^* \succ P^n$ for all n. Suppose that the former case holds. Take any $R \sim S$ and R, S compatible with P. Correlation Consistency implies for all $\alpha \in (0,1)$ and $n \geq 1$, $\alpha P^n + (1-\alpha)R \succ \alpha P^* + (1-\alpha)S$. By Continuity 1, as n goes to infinity, I have $\alpha P^* + (1-\alpha)R \succsim \alpha P^* + (1-\alpha)S$. This holds for all $R \sim S$ with R, S compatible with P. By symmetry, for all $\alpha \in (0,1)$ and $R \sim S$ with R, S compatible with P,

$$\alpha P^* + (1-\alpha)S \sim \alpha P^* + (1-\alpha)R.$$

The same result holds if $P^* \succ P^n$ for all n. Without loss of generality, I assume that $P^n \succ P^*$ for all n and $P^* \succ \tilde{P}$ from now on.

Fix any Q compatible with P and we know Q is also compatible with $\tilde{P}$, P^* and P^n for each n. By Correlation Consistency, for any $\beta \in (0,1)$, $\beta P^* + (1-\beta)Q \succ \beta\tilde{P} + (1-\beta)Q$ and $\beta P^* + (1-\beta)Q, \beta\tilde{P} + (1-\beta)Q \in M(\beta P_1 + (1-\beta)Q_1, \beta P_2 + (1-\beta)Q_2)$. Similarly, as $P^n \succ P^*$ for all n, for any $\beta \in (0,1)$, $\beta P^n + (1-\beta)Q \succ \beta P^* + (1-\beta)Q$ and $\beta P^n + (1-\beta)Q \in M(\beta P_1 + (1-\beta)Q_1, \beta P_2 + (1-\beta)Q_2)$. For any $R \sim S$ with R, S compatible with both P, Q, R, S are also compatible with $\beta P^n + (1-\beta)Q$ and $\beta P^* + (1-\beta)Q$. With the same arguments above, one can show that for any $\alpha \in (0,1)$

and $\beta \in (0,1)$,

$$\alpha[\beta P^* + (1-\beta)Q] + (1-\alpha)R \sim \alpha[\beta P^* + (1-\beta)Q] + (1-\alpha)S.$$

This can be rearranged as

$$\beta[\alpha P^* + (1-\alpha)R] + (1-\beta)[\alpha Q + (1-\alpha)R] \sim \beta[\alpha P^* + (1-\alpha)S] + (1-\beta)[\alpha Q + (1-\alpha)S].$$

Again by Continuity 1, let $\beta \to 0^+$ and I have

$$\alpha Q + (1-\alpha)R \sim \alpha Q + (1-\alpha)S, \tag{A.4}$$

for any $\alpha \in (0,1)$, $R \sim S$, Q compatible with P, R, S and P compatible with Q, R, S.

Fix $P, \tilde{P}$ and Q such that P is compatible with Q, I want to strengthen property (A.4) by discarding the constraint that R, S are compatible with P. By Continuity 3, as $P \succ \tilde{P}$, I can find $\bar{\varepsilon} > 0$ such that for all $\varepsilon = (\varepsilon^1, 0)$ with $\varepsilon^1 \in (0, \bar{\varepsilon})$, I have $P_\varepsilon \succ \tilde{P}_\varepsilon$. Note that $\tilde{P}_\varepsilon, P_\varepsilon \in M(P_{\varepsilon,1}, P_2)$. Since $\text{supp}(P_1) \cup \text{supp}(Q_1)$ is finite and P is compatible with Q, I can make $\bar{\varepsilon}$ small enough such that for all $\varepsilon^1 \in (0, \bar{\varepsilon})$, $\text{supp}(P_{\varepsilon,1}) \subset X_1^o$, $\text{supp}(P_{\varepsilon,1}) \cap \text{supp}(Q_1) = \emptyset$ and $\tilde{P}_\varepsilon, P_\varepsilon$ are compatible with Q. Then

$$\alpha Q + (1-\alpha)R \sim \alpha Q + (1-\alpha)S,$$

for any $\varepsilon \in (0, \bar{\varepsilon})$, $\alpha \in (0,1)$, $R \sim S$, Q compatible with R, S, P and P_ε compatible with R, S.

By varying ε, we can further get rid of the constraint that R, S are compatible with P_ε for some $\varepsilon \in (0, \bar{\varepsilon})$. This is guaranteed by the fact that each lottery in $\mathcal{P}$ has a finite support and $\text{supp}(P_{\varepsilon,1}) \subset X_1^o$ for each $\varepsilon^1 \in (0, \bar{\varepsilon})$. Concretely, for any $R \sim S$ with R, S compatible with Q, I can always find ε^* such that R, S are compatible with P_{ε^*}. Thus,

$$\alpha Q + (1-\alpha)R \sim \alpha Q + (1-\alpha)S,$$

for any $\alpha \in (0,1)$, $R \sim S$ and Q compatible with R, S, P. The same argument can be applied to relax the requirement that Q is compatible with P. Hence, for any $Q \in \mathcal{P}$,

$$\alpha Q + (1-\alpha)R \sim \alpha Q + (1-\alpha)S,$$

for any $\alpha \in (0,1)$, $R \sim S$, Q compatible with R, S. ■

The second main challenge is that Lemma A.2.9 only holds for product lotteries and Lemma A.3.3 does not directly imply $\alpha Q + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$ if $R \succ S$ and Q is compatible with R, S. Moreover, in Lemma A.3.3, I cannot choose $R = Q$ and hence I have not yet shown that the utility of $\alpha P + (1-\alpha)$ lies between those of P and Q even if they are compatible. My next goal is to deal with these issues.

First, for a fixed $y \in X_2$, I will focus on the set of lotteries whose utilities are strictly bounded by two lotteries in $\mathcal{L}^0(X_1) \times \{y\}$, that is,

$$\Phi_{2,y} = \Big\{P \in \mathcal{P} : \exists\, T, T' \in \mathcal{P} \text{ with } T_2 = T_2' = y \text{ s.t. } T \succ P \succ T'\Big\}.$$

Lemma A.3.4 is a counterpart of Lemma A.2.4 and part (i) of Lemma A.2.3 in $\mathcal{P}$.

Lemma A.3.4 *(i). For each $P, Q, R \in \mathcal{P}$ with $P \succ Q \succ R$, there exists $\lambda \in (0,1)$ such that $\lambda P + (1-\lambda)R \sim Q$.*
(ii). For each $P \in \mathcal{P}$, there exists $(x_1, x_2) \in X_1 \times X_2$ such that $P \sim (x_1, x_2)$. Moreover, if $P \in \Phi_{2,y}$ for some $y \in X_2$, then we can choose $x_2 = y$.

Proof of Lemma A.3.4. (i). Denote $A = \{\alpha \in (0,1) : \alpha P + (1-\alpha)R \succ Q\}$ and $\lambda = inf A$. I claim that $\lambda P + (1-\lambda)R \sim Q$. Suppose by contradiction that $\lambda P + (1-\lambda)R \not\sim Q$. If $\lambda P + (1-\lambda)R \succ Q$, then $\lambda \in A$, which is open by Continuity 1. Hence there exists $\lambda' < \lambda$ with $\lambda' \in A$, which contradicts with the definition of λ. If $\lambda P + (1-\lambda)R \prec Q$, then $\lambda \in \{\alpha \in (0,1) : \alpha P + (1-\alpha)R \prec Q\}$, which is also open. I can find $\varepsilon > 0$ such that $[\lambda, \lambda + \varepsilon) \subseteq (0,1) \backslash A$. Again a contradiction with the definition of λ. Hence $\lambda P + (1-\lambda)R \sim Q$. (Unlike part (i) of Lemma A.2.3, λ is not necessarily unique here.)

(ii). For each $P \in \mathcal{P}$, denote $x_i = \max \text{supp}(P_i), y_i = \min \text{supp}(P_i)$ for $i = 1, 2$. By Monotonicity, $(x_1, x_2) \succsim P \succsim (y_1, y_2)$. Then either $(x_1, x_2) \succsim P \succsim (x_1, y_2)$ or $(x_1, y_2) \succsim P \succsim (y_1, y_2)$. Without loss of generality, suppose the former case holds. Using the same argument in part (i) of Lemma A.2.3, there is a unique $\lambda \in [0, 1]$ such that $P \sim (x_1, \lambda\delta_{y_1} + (1-\lambda)\delta_{y_2})$. By Lemma A.2.1, there exists $x_2' \in X_2$ where $P \sim (x_1, \lambda\delta_{y_1} + (1-\lambda)\delta_{y_2}) \sim (x_1, x_2')$. If further $P \in \Phi_{2,y}$ for some $y \in X_2$, then we can find $p_1, p_1' \in \mathcal{L}^0(X_1)$ with $(p_1, y) \succ P \succ (p_1', y)$. By the same argument, there exists $x' \in X_1$ such that $P \sim (x', y)$. ■

The next lemma generalizes Lemma A.3.3 and shows that the independence axiom holds on $\Phi_{2,y}$ subject to the compatibility constraint.

Lemma A.3.5 *For each $y \in X_2$ and $P, Q, R, S \in \Phi_{2,y}$, the following properties hold:*

i). $P \sim Q$ and P compatible with $Q \Longrightarrow \alpha P + (1-\alpha)Q \sim P \sim Q$, $\forall \alpha \in (0, 1)$;

ii). $P \succ Q$ and P compatible with $Q \Longrightarrow P \succ \alpha P + (1-\alpha)Q \succ Q$, $\forall \alpha \in (0, 1)$;

iii). $P \succ Q$, $R \sim S$, P compatible with R and Q compatible with $S \Longrightarrow \alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$, $\forall \alpha \in (0, 1)$;

iv). $P \sim Q$, $R \sim S$, P compatible with R and Q compatible with $S \Longrightarrow \alpha P + (1-\alpha)R \sim \alpha Q + (1-\alpha)S$, $\forall \alpha \in (0, 1)$.

Proof of Lemma A.3.5. I will prove (i) and (ii). (iii) and (iv) can be shown similarly. Suppose $P, Q \in \Phi_{2,y}$ for some $y \in X_2$ and P, Q are compatible. By Lemma A.3.4, there exists $x_P, x_Q \in X_1^o$ such that $P \sim (x_P, y)$ and $Q \sim (x_Q, y)$. By Lemma A.2.1, there exists $\varepsilon > 0$ such that for all $z_P \in [x_P - \varepsilon, x_P], z_Q \in [x_Q - \varepsilon, x_Q]$, there exist $z_P' \geq x_P, z_Q' \geq x_Q$ such that $P \sim (1/2\delta_{z_P} + 1/2\delta_{z_P'}, y)$ and $Q \sim (1/2\delta_{z_Q} + 1/2\delta_{z_Q'}, y)$. Moreover, as z_P, z_Q increases, z_P', z_Q' will be decreasing continuously. Since P, Q are simple, that is, $\text{supp}(P_1) \cup \text{sup}(Q_1)$ is finite, I can construct $z_P^* \neq z_Q^*$, $z_P^{*'} \neq z_Q^{*'}$ and $z_P^*, z_Q^*, z_P^{*'}, z_Q^{*'} \notin \text{supp}(P_1) \cup \text{sup}(Q_1)$. Denote $P' = (1/2\delta_{z_P^*} + 1/2\delta_{z_P^{*'}}, y)$, $Q' = (1/2\delta_{z_Q^*} + 1/2\delta_{z_Q^{*'}}, y)$. Then $P \sim P', Q \sim Q'$ and P, Q, P', Q' are compatible with each other. Apply Lemma A.3.3 twice and I get for any $\alpha \in (0, 1)$,

$$\alpha P + (1-\alpha)Q \sim \alpha P + (1-\alpha)Q' \sim \alpha P' + (1-\alpha)Q'.$$

Again by Lemma A.2.1 given marginal lottery in source 2 as y,

$$P \sim Q \Longrightarrow P' \sim Q' \Longrightarrow \alpha P + (1-\alpha)Q \sim \alpha P' + (1-\alpha)Q' \sim Q' \sim Q,$$
$$P \succ Q \Longrightarrow P' \succ Q' \Longrightarrow P \sim P' \sim \alpha P + (1-\alpha)Q \sim \alpha P' + (1-\alpha)Q' \succ Q' \sim Q.$$

■

Like Lemma A.2.10, I claim that if the independence property holds on $\Phi_{2,y}$ for each y, then it also holds on their union.

Lemma A.3.6 *For each $y \in X_2$ and $P, Q, R, S \in \cup_{y \in X_2} \Phi_{2,y}$, the following hold:*

i). $P \sim Q$ and P compatible with $Q \Longrightarrow \alpha P + (1-\alpha)Q \sim P \sim Q$, $\forall \alpha \in (0,1)$;

ii). $P \succ Q$ and P compatible with $Q \Longrightarrow P \succ \alpha P + (1-\alpha)Q \succ Q$, $\forall \alpha \in (0,1)$;

iii). $P \succ Q$, $R \sim S$, P compatible with R and Q compatible with $S \Longrightarrow \alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$, $\forall \alpha \in (0,1)$;

iv). $P \sim Q$, $R \sim S$, P compatible with R and Q compatible with $S \Longrightarrow \alpha P + (1-\alpha)R \sim \alpha Q + (1-\alpha)S$, $\forall \alpha \in (0,1)$.

Proof of Lemma A.3.6. First, for any $P, Q, R, S \in \cup_{y \in X_2} \Phi_{2,y}$, I claim that there exist finitely many $z_k \in X_2, k = 1, ..., K$ such that $z_1 < z_2 < ... < z_K$ and $P, Q, R, S \in \cup_{k=1}^{K} \Phi_{2,z_k}$. Choose $P^1, P^2 \in \{P, Q, R, S\}$ such that $P^1 \succsim P, Q, R, S \succsim P^2$. Suppose that $P^1 \in \Phi_{2,y_1}$ and $P^2 \in \Phi_{2,y_2}$ with $y_1 \geq y_2$. If $y_1 = y_2$, then $P, Q, R, S \in \Phi_{2,y_1}$.

Now suppose that $y_1 > y_2$ and by Lemma A.2.1, I can find $t, t' \in \mathcal{L}^0(X_1)$ with $(t, y_1) \succ P^1 \succ (t', y_1), (t, y_2) \succ P^2 \succ (t', y_2)$. Notice that for each $y \in [y_2, y_1]$, $H(y) := \{P \in \mathcal{P} : (t, y) \succ P \succ (t', y)\} \subseteq \Phi_{2,y}$. By Continuity, $\{P \in \mathcal{P} : P^1 \succsim P \succsim P^2\} \subseteq \cup_{y_2 \leq y \leq y_1} H(y)$ and for all $y \in [y_2, y_1]$, there exists $\varepsilon_y > 0$ such that $H(y) \cap H(y') \neq \emptyset$ for all $y' \in [y - \varepsilon_y, y + \varepsilon_y] \cap [y_2, y_1]$. By Finite Cover Theorem, I can find finitely many $z_1 < z_2 < ... < z_K \in [y_2, y_1]$ with $z_1 = y_2, z_K = y_1$ and $[y_2, y_1] \subseteq \cup_{k=1}^{K} [z_k - \varepsilon_{z_k}, z_k + \varepsilon_{z_k}]$. This implies

$$P, Q, R, S \in \{P \in \mathcal{P} : P^1 \succsim P \succsim P^2\} \subseteq \cup_{y_2 \leq y \leq y_1} H(y) = \cup_{k=1}^{K} H(z_k) \subseteq \cup_{k=1}^{K} \Phi_{2,z_k}.$$

Then I will use induction to show that the four properties hold for $P, Q, R, S \in \cup_{k=1}^{K} \Phi_{2,z_k}$. The proof idea is similar to the proof of Lemma A.2.10. By Lemma A.3.5, the four properties hold if $P, Q, R, S \in \Phi_{2,z_1}$. Suppose by induction that they also hold if $P, Q, R, S \in \cup_{k=1}^{t} \Phi_{2,z_k}$ for some $1 \leq t < K$. By my construction of $\{z_k\}$, $\Phi_{2,z_t} \cap \Phi_{2,z_{t+1}}$ is nonempty and I can find $T^1, T^2 \in \Phi_{2,z_t} \cap \Phi_{2,z_{t+1}}$ with $T^1 \succ T^2$. By Lemma A.3.4 and Lemma A.2.1, since P_1, Q_1, R_1, S_1 have finite supports, I can also construct $p_1, p_2, q_1, q_2 \in \mathcal{L}^0(X_1)$ such that $(p_1, z_{t+1}) \sim (q_1, z_t) \sim T^1$, $(p_2, z_{t+1}) \sim (q_2, z_t) \sim T^2$ and $(p_1, z_{t+1}), (q_1, z_t), (p_2, z_{t+1}), (q_2, z_t)$ are compatible with P, Q, R, S.

For property (i), suppose $P \succsim Q$, P is compatible with Q and $P, Q \in \cup_{k=1}^{t+1} \Phi_{2,z_k}$. Since $P \sim Q$, $P, Q \in \Phi_{2,z_k}$ for some $k = 1, ..., t+1$ and hence (i) holds.

For property (ii), if $P \succ Q$, then it suffices to consider the case where $P \in \Phi_{2,z_{t+1}} \backslash (\cup_{k=1}^{t} \Phi_{2,z_k})$ and $Q \in (\cup_{k=1}^{t} \Phi_{2,z_k}) \backslash \Phi_{2,z_{t+1}}$. This implies $P \succ T^1 \succ T^2 \succ Q$. By Lemma A.3.4, there exists $\lambda_1 \neq \lambda_2 \in (0,1)$ such that $T^1 \sim \lambda_1 P + (1-\lambda_1) Q$ and $T^2 \sim \lambda_2 P + (1-\lambda_2) Q$. Then property (ii) holds for $\lambda = \lambda_1, \lambda_2$.

Suppose that $\lambda_i > \lambda_{-i}$ for some $i = 1, 2$. By Lemma A.3.4, we can find $P', Q' \in \hat{\mathcal{P}}$ with $Q' \sim Q, P' \sim P$ and $P, P', Q, Q', (p_1, z_{t+1}), (q_1, z_t), (p_2, z_{t+1}), (q_2, z_t)$ compatible with each other. This guarantees

$$T^1 \sim \lambda_1 P + (1-\lambda_1) Q \sim \lambda_1 P' + (1-\lambda_1) Q \sim \lambda_1 P + (1-\lambda_1) Q' \sim \lambda_1 P' + (1-\lambda_1) Q',$$
$$T^2 \sim \lambda_2 P + (1-\lambda_2) Q \sim \lambda_2 P' + (1-\lambda_2) Q \sim \lambda_2 P + (1-\lambda_2) Q' \sim \lambda_2 P' + (1-\lambda_2) Q'.$$

By property (i), for all $\beta, \beta' \in (0,1)$, $\beta P + (1-\beta) P' \sim P, \beta' Q + (1-\beta') Q' \sim Q$. Apply Lemma A.3.3 twice and then for each $\lambda, \beta, \beta' \in (0,1)$

$$\lambda P + (1-\lambda) Q \sim \lambda(\beta P + (1-\beta) P') + (1-\lambda)(\beta' Q + (1-\beta') Q'). \tag{A.5}$$

For any $\lambda \in (\lambda_{-i}, \lambda_i)$, let $\beta = 1$, $\beta' = \frac{\lambda_i - \lambda}{\lambda_i (1-\lambda)}$, and (A.5) becomes

$$\begin{aligned} \lambda P + (1-\lambda) Q \sim & \frac{\lambda}{\lambda_i} (\lambda_i P + (1-\lambda_i) Q') + (1 - \frac{\lambda}{\lambda_i}) Q \\ \sim & \frac{\lambda}{\lambda_i} (q_i, z_t) + (1 - \frac{\lambda}{\lambda_i}) Q \end{aligned}$$

The second indifference comes from the fact that $\lambda_i P + (1-\lambda_i)Q' \sim T^i \sim (q_i, z_t)$ and Lemma A.3.3. Then by the inductive hypothesis on $\cup_{k=1}^{t}\Phi_{2,z_k}$, I have

$$P \succ (q_i, z_t) \succ \lambda P + (1-\lambda)Q \sim \frac{\lambda}{\lambda_i}(q_i, z_t) + (1-\frac{\lambda}{\lambda_i})Q \succ Q.$$

If $\lambda > \lambda_i$, then let $\beta = \frac{\lambda-\lambda_i}{\lambda(1-\lambda_i)}$, $\beta' = 0$ and (A.5) becomes

$$\begin{aligned}\lambda P + (1-\lambda)Q \sim & \frac{\lambda-\lambda_i}{1-\lambda_i}P + (1-\frac{\lambda-\lambda_i}{1-\lambda_i})(\lambda_i P' + (1-\lambda_i)Q) \\ \sim & \frac{\lambda-\lambda_i}{1-\lambda_i}P + (1-\frac{\lambda-\lambda_i}{1-\lambda_i})(p_i, z_{t+1}).\end{aligned}$$

The second indifference comes from the fact that $\lambda_i P' + (1-\lambda_i)Q \sim T^i \sim (p_i, z_{t+1})$ and Lemma A.3.3. Then by Lemma A.3.5 on $\Phi_{2,z_{t+1}}$, I have

$$P \succ \lambda P + (1-\lambda)Q \sim \frac{\lambda-\lambda_i}{1-\lambda_i}P + (1-\frac{\lambda-\lambda_i}{1-\lambda_i})(p_i, z_{t+1}) \succ (p_i, z_{t+1}) \succ Q.$$

A symmetric argument works for $\lambda < \lambda_{-i}$. Hence property (ii) holds on $\cup_{k=1}^{t+1}\Phi_{2,z_k}$.

I claim that for $P, Q \in \cup_{k=1}^{t+1}\Phi_{2,z_k}$ with $P \succ Q$, P compatible with Q and $1 > \lambda_1 > \lambda_2 > 0$, we have $\lambda_1 P + (1-\lambda_1)Q \succ \lambda_2 P + (1-\lambda_2)Q$. To see this, by (A.5), I can find $P' \sim P$ where P' is compatible with both P and Q such that

$$\lambda_1 P + (1-\lambda_1)Q \sim \frac{\lambda_1-\lambda_2}{1-\lambda_2}P' + \frac{1-\lambda_1}{1-\lambda_2}[\lambda_2 P + (1-\lambda_2)Q] \succ \lambda_2 P + (1-\lambda_2)Q.$$

The second strict ranking follows from property (ii) since $P \sim P' \succ \lambda_2 P + (1-\lambda_2)Q$.

Given this claim, the proof of (iii) and (iv) on $\cup_{k=1}^{t+1}\Phi_{2,z_k}$ is similar to the proof of (ii). By induction, the four properties hold for $P, Q, R, S \in \cup_{k=1}^{K}\Phi_{2,z_k}$ and hence arbitrary $P, Q, R, S \in \cup_{y\in X_2}\Phi_{2,y}$. ■

It is worthwhile to note that $\cup_{y\in X_2}\Phi_{2,y}$ is not the same as $\mathcal{P}$. The next lemmas shows that they only differ in the worst and the best possible lotteries.

Lemma A.3.7 $\mathcal{P}\backslash(\cup_{y\in X_2}\Phi_{2,y}) = \{(\underline{c}_1, \underline{c}_2), (\overline{c}_1, \overline{c}_2)\}$.

Proof of Lemma A.3.7. First, for each $P \in \mathcal{P}$ with $(\overline{c}_1, \overline{c}_2) \succ P \succ (\underline{c}_1, \underline{c}_2)$, there exists $Q, Q' \in \cup_{y \in X_2} \Phi_{2,y}$ with $Q \succ P \succ Q'$, which implies $Q \in \cup_{y \in X_2} \Phi_{2,y}$. Hence

$$\mathcal{P} = (\cup_{y \in X_2} \Phi_{2,y}) \cup \{P \in \mathcal{P} : P \sim (\underline{c}_1, \underline{c}_2) \text{ or } P \sim (\overline{c}_1, \overline{c}_2)\}.$$

Then it suffices to show that $P \sim (\underline{c}_1, \underline{c}_2)$ if and only if $P = (\underline{c}_1, \underline{c}_2)$, and $P \sim (\overline{c}_1, \overline{c}_2)$ if and only if $P = (\overline{c}_1, \overline{c}_2)$. This is trivial by Axiom M as for any $P \neq (\underline{c}_1, \underline{c}_2), (\overline{c}_1, \overline{c}_2)$, P dominates $(\underline{c}_1, \underline{c}_2)$ and is dominated by $(\overline{c}_1, \overline{c}_2)$. ■

Since $P \succ (\underline{c}_1, \underline{c}_2)$ for any $P \neq (\underline{c}_1, \underline{c}_2)$ and $P \prec (\overline{c}_1, \overline{c}_2)$ for any $P \neq (\overline{c}_1, \overline{c}_2)$, one can easily use the arguments in Lemma A.3.5 to show that the independence property holds for $P, Q, R, S \in \Phi_{2,0} \cup \{(\underline{c}_1, \underline{c}_2)\}$ or $P, Q, R, S \in \Phi_{2,\overline{c}_2} \cup \{(\overline{c}_1, \overline{c}_2)\}$. This leads to the following corollary of Lemma A.3.6.

Corollary A.3.8 *The four properties in Lemma A.3.6 hold on $\mathcal{P}$.*

Finally I slightly relax the requirement of compatibility. For each $P, Q \in \mathcal{P}$, I say P and Q are *weakly compatible* if the they satisfy the following conditions:

- $\text{supp}(P_1) \cap \text{supp}(Q_1) \subseteq \{\underline{c}_1, \overline{c}_1\}$;
- If $\underline{c}_1 \in \text{supp}(P_1) \cap \text{supp}(Q_1)$, then $P_{2|\underline{c}_1} = Q_{2|\underline{c}_1} = \underline{c}_2$;
- If $\overline{c}_1 \in \text{supp}(P_1) \cap \text{supp}(Q_1)$, then $P_{2|\overline{c}_1} = Q_{2|\overline{c}_1} = \overline{c}_2$.

In other words, if P is weakly compatible with Q, then both P_1 and P_2 attach positive probability to $\underline{c}_1$ (or $\overline{c}_1$) only if the conditional lottery in source 2 given $\underline{c}_1$ (or $\overline{c}_1$) is degenerate at $\underline{c}_2$ (or $\overline{c}_2$) for both P and Q.

Lemma A.3.9 *The four properties in Lemma A.3.6 hold on $\mathcal{P}$ if "compatible" is replaced with "weakly compatible".*

Proof of Lemma A.3.9. Suppose P, Q are weakly compatible but not compatible, that is, $\emptyset \neq \text{supp}(P_1) \cap \text{supp}(Q_1) \subseteq \{\underline{c}_1, \overline{c}_1\}$. I claim that I can find $\tilde{P} \sim P$ and $\tilde{Q} \sim Q$ such that $\tilde{P}$ is compatible with $\tilde{Q}$ and for any $\alpha \in (0, 1)$, $\alpha P + (1-\alpha)Q \sim \alpha\tilde{P} + (1-\alpha)\tilde{Q}$, unless $P = Q = (\underline{c}_1, \underline{c}_2)$ or $P = Q = (\overline{c}_1, \overline{c}_2)$.

Case 1: If $\text{supp}(P_1) \cap \text{supp}(Q_1) = \{\underline{c}_1, \overline{c}_1\}$, then $P_{2|\underline{c}_1} = Q_{2|\underline{c}_1} = \underline{c}_2$ and $P_{2|\overline{c}_1} = Q_{2|\overline{c}_1} = \overline{c}_2$. First, I suppose that $P_1(\underline{c}_1)+P_1(\overline{c}_1) < 1$ or $Q_1(\underline{c}_1)+Q_1(\overline{c}_1) < 1$. By symmetry, it suffices to focus on the former case. Denote $P^o = \sum_{x \neq \underline{c}_1, \overline{c}_1} \frac{P_1(x)}{1-P_1(\underline{c}_1)-P_1(\overline{c}_1)}(x, P_{2|x})$.

$$P = P_1(\underline{c}_1)(\underline{c}_1, \underline{c}_2) + P_1(\overline{c}_1)(\overline{c}_1, \overline{c}_2) + (1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))P^o.$$

P^o and Q are compatible and $\underline{c}_1, \overline{c}_1 \notin \text{supp}(P_1^o)$. I can similarly define Q^o if $Q_1(\underline{c}_1) + Q_1(\overline{c}_1) < 1$. Otherwise, just choose an arbitrary Q^o so long as $\underline{c}_1, \overline{c}_1 \notin \text{supp}(Q_1^o)$.

By Monotonicity, $(\overline{c}_1, \overline{c}_2) \succ P^o \succ (\underline{c}_1, \underline{c}_2)$. Then one can find $\varepsilon_P > 0$ such that $\overline{c}_1 - \varepsilon_P, \varepsilon_P \notin \text{supp}(P_1) \cup \text{supp}(Q_1)$, $\varepsilon_P \neq \overline{c}_1 - \varepsilon_P$ and $(\overline{c}_1 - \varepsilon_P, \overline{c}_2) \succ P^o \succ (\underline{c}_1 + \varepsilon_P, \underline{c}_2)$.

By Lemma A.3.4, there exists $\lambda_P \in (0, 1)$ such that

$$P^o \sim \lambda_P(\overline{c}_1 - \varepsilon_P, \overline{c}_2) + (1 - \lambda_P)(\underline{c}_1 + \varepsilon_P, \underline{c}_2) := P^{o\prime},$$

and $P^{o\prime}$ is compatible with P, Q. By Corollary A.3.8,

$$P' := P_1(\underline{c}_1)(\underline{c}_1, \underline{c}_2) + P_1(\overline{c}_1)(\overline{c}_1, \overline{c}_2) + (1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))P^{o\prime} \sim P.$$

Notice that

$$\begin{aligned} P' =& [P_1(\underline{c}_1)(\underline{c}_1, \underline{c}_2) + (1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))(1 - \lambda_P)(\underline{c}_1 + \varepsilon_P, \underline{c}_2)] \\ &+ [P_1(\overline{c}_1)(\overline{c}_1, \overline{c}_2) + (1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))\lambda_P(\overline{c}_1 - \varepsilon_P, \overline{c}_2)]. \end{aligned}$$

By Lemma A.2.1 given $\underline{c}_2$ or $\overline{c}_2$ in source two, there exists $\overline{p}, \underline{p} \in \mathcal{L}^0(X_1)$ with $(\overline{p}, \overline{c}_2), (\underline{p}, \underline{c}_2), P, P', Q, Q'$ are pairwise compatible and

$$\begin{aligned} (\overline{p}, \overline{c}_2) &\sim \frac{P_1(\overline{c}_1)(\overline{c}_1, \overline{c}_2) + \lambda_P(1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))(\overline{c}_1 - \varepsilon_P, \overline{c}_2)}{P_1(\overline{c}_1) + \lambda_P(1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))}; \\ (\underline{p}, \underline{c}_2) &\sim \frac{P_1(\underline{c}_1)(\underline{c}_1, \underline{c}_2) + (1 - \lambda_P)(1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))(\underline{c}_1 + \varepsilon_P, \underline{c}_2)}{P_1(\underline{c}_1) + (1 - \lambda_P)(1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))}. \end{aligned}$$

It is important to observe that $\overline{p} \neq \overline{c}_1$ and $\underline{p} \neq \underline{c}_1$.

Again by Corollary A.3.8,

$$\tilde{P} := \lambda_P^*(\underline{p}, \underline{c}_2) + (1 - \lambda_P^*)(\overline{p}, \overline{c}_2) \sim P' \sim P,$$

where $\lambda_P^* = P_1(\underline{c}_1) + (1 - \lambda_P)(1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))$. $\tilde{P}$ is compatible with P and Q.

I claim that for each $\alpha \in (0,1)$, $\alpha P + (1-\alpha)Q \sim \alpha\tilde{P} + (1-\alpha)Q$. To see this, notice that for each $\alpha \in (0,1)$,

$$\begin{aligned}\alpha P + (1-\alpha)Q = &(\alpha P_1(\underline{c}_1) + (1-\alpha)Q_1(\underline{c}_1))(\underline{c}_1, \underline{c}_2) + (\alpha P_1(\overline{c}_1) + (1-\alpha)Q_1(\overline{c}_1))(\overline{c}_1, \overline{c}_2) \\ &+ \alpha(1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))P^o + (1-\alpha)(1 - Q_1(\underline{c}_1) - Q_1(\overline{c}_1))Q^o.\end{aligned}$$

Since $P^o \sim P^{o\prime}$, P^o and $P^{o\prime}$ are compatible with Q and P, Corollary A.3.8 implies

$$\begin{aligned}\alpha P + (1-\alpha)Q \sim &(\alpha P_1(\underline{c}_1) + (1-\alpha)Q_1(\underline{c}_1))(\underline{c}_1, \underline{c}_2) + (\alpha P_1(\overline{c}_1) + (1-\alpha)Q_1(\overline{c}_1))(\overline{c}_1, \overline{c}_2) \\ &+ \alpha(1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))P^{o\prime} + (1-\alpha)(1 - Q_1(\underline{c}_1) - Q_1(\overline{c}_1))Q^o \\ = &\alpha P' + (1-\alpha)Q \\ = &\alpha[P_1(\underline{c}_1)(\underline{c}_1, \underline{c}_2) + (1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))(1 - \lambda_P)(\underline{c}_1 + \varepsilon_P, \underline{c}_2)] \\ &+ (1-\alpha)Q_1(\underline{c}_1)(\underline{c}_1, \underline{c}_2) \\ &+ \alpha[P_1(\overline{c}_1)(\overline{c}_1, \overline{c}_2) + (1 - P_1(\underline{c}_1) - P_1(\overline{c}_1))\lambda_P(\overline{c}_1 - \varepsilon_P, \overline{c}_2)] \\ &+ (1-\alpha)Q_1(\overline{c}_1)(\overline{c}_1, \overline{c}_2) \\ &+ (1-\alpha)(1 - Q_1(\underline{c}_1) - Q_1(\overline{c}_1))Q^o.\end{aligned}$$

The first two terms in the last expression have $\underline{c}_2$ in source 2, while the third and the fourth terms have $\overline{c}_2$ in the source 2. Apply Lemma A.2.1 given $\underline{c}_2$ or $\overline{c}_2$ in source 2 and Corollary A.3.8 sequentially, we can derive

$$\begin{aligned}\alpha P + (1-\alpha)Q \sim &\alpha\lambda_P^*(\underline{p}, \underline{c}_2) + (1-\alpha)Q_1(\underline{c}_1)(\underline{c}_1, \underline{c}_2) \\ &+ \alpha(1 - \lambda_P^*)(\overline{p}, \overline{c}_2) + (1-\alpha)Q_1(\overline{c}_1)(\overline{c}_1, \overline{c}_2) \\ &+ (1-\alpha)(1 - Q_1(\underline{c}_1) - Q_1(\overline{c}_1))Q^o \\ = &\alpha\tilde{P} + (1-\alpha)Q.\end{aligned}$$

Now suppose $P_1(\underline{c}_1) + P_1(\overline{c}_1) = Q_1(\underline{c}_1) + Q_1(\overline{c}_1) = 1$. As $\text{supp}(P_1) \cap \text{supp}(Q_1) = \{\underline{c}_1, \overline{c}_1\}$, $P_1(\underline{c}_1), Q_1(\underline{c}_1) \in (0,1)$. By Monotonicity, $(\overline{c}_1, \overline{c}_2) \succ P, Q \succ (\overline{c}_1, \overline{c}_2)$. Then there is $\tilde{P} \sim P$ such that $\underline{c}_1, \overline{c}_1 \notin \text{supp}(\tilde{P}_1)$ and $\tilde{P}$ is compatible with P and Q. For any $\beta \in (0,1)$, by Corollary A.3.8, $P \sim \beta P + (1-\beta)P' := P^\beta$. Clearly, $P^\beta(\underline{c}_1) + P^\beta(\overline{c}_1) < 1$. By applying the previous result to P^β and Q, for each β, there exists $\tilde{P}^\beta \sim P^\beta \sim P$ with $\tilde{P}^\beta$ compatible with Q and for each $\alpha \in (0,1)$, $\alpha P^\beta + (1-\alpha)Q \sim \alpha \tilde{P}^\beta + (1-\alpha)Q$. Again by Corollary A.3.8, I can actually choose $\tilde{P}^\beta$ to be the same across all $\beta \in (0,1)$. Denote it by $\tilde{P}$. Hence, for each $\beta, \alpha \in (0,1)$,

$$\alpha\beta P + \alpha(1-\beta)P' + (1-\alpha)Q \sim \alpha\tilde{P} + (1-\alpha)Q$$

By Continuity 1, let $\beta \to 1$ and then for each $\alpha \in (0,1)$, $\alpha P + (1-\alpha)Q \sim \alpha\tilde{P} + (1-\alpha)Q$.

Case 2: If $\text{supp}(P_1) \cap \text{supp}(Q_1) = \{\underline{c}_1\}$ or $\{\overline{c}_1\}$, then similar to Case 1.

As a summary, either $P = Q = (\underline{c}_1, \underline{c}_2)$ or $P = Q = (\overline{c}_1, \overline{c}_2)$ or there exists $\tilde{P} \sim P$ and $\tilde{Q} \sim Q$ such that $\tilde{P}$ is compatible with $\tilde{Q}$ and for any $\alpha \in (0,1)$, $\alpha P + (1-\alpha)Q \sim \alpha\tilde{P} + (1-\alpha)\tilde{Q}$. Now I am ready to prove the four properties.

For (i) and (ii), if $P = Q = (\underline{c}_1, \underline{c}_2)$ or $P = Q = (\overline{c}_1, \overline{c}_2)$, then the result is trivial. Otherwise, there exists $\tilde{P} \sim P$ and $\tilde{Q} \sim Q$ such that for any $\alpha \in (0,1)$,

$$P \sim Q \Longrightarrow P' \sim Q' \Longrightarrow \alpha P + (1-\alpha)Q \sim \alpha P' + (1-\alpha)Q' \sim P' \sim P,$$
$$P \succ Q \Longrightarrow P' \succ Q' \Longrightarrow P \sim P' \succ \alpha P + (1-\alpha)Q \sim \alpha P' + (1-\alpha)Q' \succ Q' \sim Q.$$

For (iii) and (iv), if $P = R = (\underline{c}_1, \underline{c}_2)$ or $Q = S = (\underline{c}_1, \underline{c}_2)$ or $P = R = (\overline{c}_1, \overline{c}_2)$ or $Q = S = (\overline{c}_1, \overline{c}_2)$, then by Monotonicity, the primitives of (iii) or (iv) hold only if $P = Q = R = S$, or $P = R = S = (\overline{c}_1, \overline{c}_2) \succ Q$ or $P \succ R = S = Q = (\underline{c}_1, \underline{c}_2)$, in which cases the result holds trivially. By excluding those cases, we can construct $\tilde{P} \sim P$, $\tilde{Q} \sim Q$, $\tilde{R} \sim R$ and $\tilde{S} \sim S$ such that $\tilde{P}$ is compatible with $\tilde{R}$, $\tilde{Q}$ is compatible with $\tilde{S}$ and for any $\alpha \in (0,1)$, $\alpha P + (1-\alpha)R \sim \alpha\tilde{P} + (1-\alpha)\tilde{R}$, $\alpha Q + (1-\alpha)S \sim \alpha\tilde{Q} + (1-\alpha)\tilde{S}$. By Corollary A.3.8,

$$P \sim Q, R \sim S \Longrightarrow \alpha P + (1-\alpha)R \sim \alpha\tilde{P} + (1-\alpha)\tilde{R} \sim \alpha\tilde{Q} + (1-\alpha)\tilde{S} \sim \alpha Q + (1-\alpha)S,$$

$$P \succ Q, R \sim S \Longrightarrow \alpha P + (1-\alpha)R \sim \alpha \tilde{P} + (1-\alpha)\tilde{R} \succ \alpha \tilde{Q} + (1-\alpha)\tilde{S} \sim \alpha Q + (1-\alpha)S.$$

This completes the proof. ■

Now I can state the main intermediate result in this proof.

Lemma A.3.10 *$\succsim$ is represented by U where for each $P \in \mathcal{P}$,*

$$U(P) = \sum_x \hat{u}_D(x, CE_{v_x}(P_{2|x}))P_1(x) \tag{A.6}$$

with continuous and strictly monotonic $\hat{u}_D$ and v_x for each $x \in X_1$. $\hat{u}_D$ and v_x for each $x \in X_1$ are unique up to positive affine transformations.

The utility function (A.6) differs from V^{ANB} since the preference for conditional lotteries in source 2 is represented by v_x, which can depend on the outcome x in source 1. If v_x is independent of x, then (A.6) agrees with V^{ANB}. If v_x is a positive affine transformation of $\hat{u}_D(x, \cdot)$, then (A.6) agrees with V^{EU}.

Proof of Lemma A.3.10. The proof can be decomposed into the following steps. First, I construct a mixture space $\mathcal{D}^*$ which is richer than $\mathcal{P}$ and argue that $\mathcal{P}$ can be mapped bijectively to a subset of the $\mathcal{D}^*$. Second, I construct an extension of $\succsim$ from $\mathcal{P}$ to $\succsim^*$ the richer domain $\mathcal{D}^*$. Third, I prove that $\succsim^*$ satisfies the standard EU axioms on $\mathcal{D}^*$ and hence is represented by some EU utility function. Finally, I show that the restriction of the utility function on $\mathcal{P}$ can be written as (A.6).

Step 1. Denote $\mathcal{D}^* := \mathcal{L}^0(X_1 \times \mathcal{L}^0(X_2))$. $\mathcal{D}^*$ is a mixture space over $X_1 \times \mathcal{L}^0(X_2)$. Following Kreps and Porteus (1978), I call each element $d \in \mathcal{D}^*$ a *temporal lottery*, which involves temporal resolution of uncertainty. I also define

$$\hat{\mathcal{D}}^* := \Big\{ d \in \mathcal{D}^* : d(x,p)d(x,p') = 0, \forall x \in X_1, p \neq p' \in \mathcal{L}^0(X_2) \Big\}.$$

$\hat{\mathcal{D}}^*$ is the set of temporal lotteries where there is no early resolution of uncertainty.

Note that for $i = 1, 2$, X_i with the standard relative topology in $\mathbb{R}$ is separable. By Kreps and Porteus (1978), $\mathcal{P}$ and $\mathcal{D}^*$ with weak topology can be metrizable by

the Prokhorov metric. Endow $\hat{\mathcal{P}}$ and $\hat{\mathcal{D}}^*$ with the relative topology with respect to the weak topology on $\mathcal{P}$ and $\mathcal{D}^*$ respectively.

Define a mapping $f : \mathcal{P} \to \hat{\mathcal{D}}^*$ as follows: for $P \in \mathcal{P}$, denote $f[P] = d \in \mathcal{D}^*$ such that for any $(x, q) \in X_1 \times \mathcal{L}^0(X_2)$, $f[P](x, q) = P_1(x)$ if $q = P_{2|x}$ and $f[P](x, q) = 0$ if $q \neq P_{2|x}$. Clearly $f[P] \in \hat{\mathcal{D}}^*$ and f is well-defined. Inversely, $f^{-1} : \hat{\mathcal{D}}^* \to \mathcal{P}$ such that $f^{-1}[d](x, y) = \sum_{q \in \mathcal{L}^0(X_2)} d(x, q)q(y)$. This is also well-defined since for each $x \in X_1$ there exists at most one $q \in \mathcal{L}^0(X_2)$ with $d(x, q) > 0$ for any $d \in \hat{\mathcal{D}}^*$. Thus, f is a bijective mapping between $\mathcal{P}$ and $\hat{\mathcal{D}}^*$. It is worth noting that f is not a homeomorphism as f is not continuous, although f^{-1} is continuous.

I define a binary relation $\succsim'$ on $\hat{\mathcal{D}}^*$ such that $d \succsim' d'$ if and only if $f^{-1}[d] \succsim f^{-1}(d')$. $\succ'$ and $\sim'$ are defined correspondingly. The next result follows from Lemma A.3.9.

Corollary A.3.11 *Suppose $\lambda_i > 0$ for all i and $\sum_{i=1}^n \lambda_i = 1$. If $d^1 = \sum_{i=1}^n \lambda_i \delta_{(x_i,p_i)}$, $d^2 = \sum_{i=1}^n \lambda_i \delta_{(y_i,q_i)}$ with $x_i \neq x_j, y_i \neq y_j$ for all $i \neq j$ and $(x_i, p_i) \sim (y_i, q_i)$ for all i, then $d^1 \sim' d^2$.*

Proof of Corollary A.3.11. By Lemma A.3.9, as $x_i \neq x_j, y_i \neq y_j$ for all $i \neq j$ and $(x_i, p_i) \sim (y_i, q_i)$ for all i, I have

$$\frac{\lambda_1}{\lambda_1 + \lambda_2}(x_1, p_1) + \frac{\lambda_2}{\lambda_1 + \lambda_2}(x_2, p_2) \sim \frac{\lambda_1}{\lambda_1 + \lambda_2}(y_1, q_1) + \frac{\lambda_2}{\lambda_1 + \lambda_2}(y_2, q_2).$$

Then by induction, I can get

$$\sum_{i=1}^n \lambda_i(x_i, p_i) \sim \sum_{i=1}^n \lambda_i(y_i, q_i).$$

Since $d^1 = f^{-1}(\sum_{i=1}^n \lambda_i(x_i, p_i))$, $d^2 = f^{-1}(\sum_{i=1}^n \lambda_i(y_i, q_i))$, I conclude that $d^1 \sim' d^2$. ■

Step 2. Then I want to extend $\succsim'$ on $\hat{\mathcal{D}}^*$ to $\succsim^*$ on $\mathcal{D}^*$. For any $d \in \mathcal{D}^* \backslash \hat{\mathcal{D}}^*$, denote by $d_1 \in \mathcal{L}^0(X_1)$ the marginal lottery in source 1 and by $d_{2|x} \in \mathcal{L}^0(\mathcal{L}^0(X_2))$ the lottery over marginal lotteries conditional on outcome x in source 1. Denote $\text{supp}(d_1) = \{x_1, ..., x_N\} \subset X_1$ with $x_1 < \cdots < x_N$ and for each $k = 1, ..., N$, $\text{supp}(d_{2|x_k}) = \{p_{k,1}, ..., p_{k,t_k}\} \subset \mathcal{L}^0(X_2)$ with $t_k \geq 1$. Since $d \notin \hat{\mathcal{D}}^*$, there exists some k' with $t_{k'} > 1$. I will map d to some temporal lottery in $\hat{\mathcal{D}}^*$ using the following algorithm.

Stage 1. First focus on outcome x_1 in source 1. Denote $z_{1,1} = x_1$. If $t_1 = 1$, by Lemma A.2.1, I can find $\hat{z}_{1,1}$ with $(z_{1,1}, \hat{z}_{1,1}) \sim (z_{1,1}, p_{1,1})$. Define $d^1 \in \mathcal{D}^*$ such that $d^1(z_{1,1}, \hat{z}_{1,1}) = d(z_{1,1}, p_{1,1})$, $d^1(z_{1,1}, q) = 0$ for all $q \neq \hat{z}_{1,1}$ and $d^1(x, q) = d(x, q)$ for all $x \neq z_{1,1}$ and $q \in \mathcal{L}^0(X_2)$.

If $t_1 > 1$ and $d(\underline{c}_1, \underline{c}_2) > 0$, that is, $x_1 = \underline{c}_1$ and $\underline{c}_2 \in \text{supp}(d_{2|\underline{c}_1})$, then I reorder lotteries in $\text{supp}(d_{2|\underline{c}_1})$ such that $(x_1, p_{1,1}) \precsim (x_1, p_{1,2}) \precsim \cdots \precsim (x_1, p_{1,t_1})$. By Monotonicity, $(x_1, p_{1,1}) = (\underline{c}_1, \underline{c}_2) \prec (x_1, p_{1,i})$ for each $i > 1$. By Continuity 3, there exists $\overline{z}_1 \in (x_1, \frac{x_1+x_2}{2})$ with $(x_1, p_{1,2}) \succ (\overline{z}_1, \underline{c}_2)$. Define $z_{1,j} = z_{1,1} + \frac{j-1}{t_1}(\overline{z}_1 - z_{1,1})$ for all $j = 1, ..., t_1$. Clearly, $z_{1,1} = x_1 < z_{1,2} < \cdots < z_{1,t_1} < \overline{z}_1$. By Monotonicity and Lemma A.2.1, I can find $\hat{z}_{1,j}$ for $j \geq 1$ with $(z_{1,j}, \hat{z}_{1,j}) \sim (x_1, p_{1,j})$. Then I define $d^1 \in \mathcal{D}^*$ such that for any $x \notin \{z_{1,j}\}_{j=1}^{t_1}$, $q \in \mathcal{L}^0(X_2)$, $d^1(x, q) = d(x, q)$, and for each $j = 1, ..., t_1$, $d^1(z_{1,j}, \hat{z}_{1,j}) = d(x_1, p_{1,j})$ and $d^1(z_{1,j}, q) = 0$ for $q \neq \hat{z}_{1,j}$.

If $t_1 > 1$ and $d(\underline{c}_1, \underline{c}_2) = 0$, then I can construct d^1 similarly by choosing $z_{1,1} > \underline{c}_1$.

Stage $i \geq 2$. Suppose that I have constructed d^{i-1}. Now consider x_i in source 1 with $x_i > x_{i-1} \geq x_1 \geq \underline{c}_1$. Denote $z_{i,1} = x_i$. If $t_i = 1$, then by Lemma A.2.1, there exists $\hat{z}_{i,1}$ with $(z_{i,1}, \hat{z}_{i,1}) \sim (z_{i,1}, p_{i,1})$. Define $d^i \in \mathcal{D}^*$ such that $d^i(z_{i,1}, \hat{z}_{i,1}) = d(z_{i,1}, p_{i,1})$, $d^i(z_{i,1}, q) = 0$ for all $q \neq \hat{z}_{i,1}$ and $d^i(x, q) = d^{i-1}(x, q)$ for all $x \neq z_{i,1}$ and $q \in \mathcal{L}^0(X_2)$. If instead $t_i > 1$, again we can reorder the marginal lotteries so that $(x_i, p_{i,1}) \precsim \cdots \precsim (x_i, p_{i,t_i})$. As $x_i > x_1 \geq \underline{c}_1$, $(x_i, p_{i,1}) \succ (\frac{x_i+x_{i-1}}{2}, \underline{c}_2) \succ (x_1, \underline{c}_2)$ by Monotonicity. Since $p_{i,1} \neq p_{i,2}$, $(x_i, p_{i,2}) \succ (x_i, \underline{c}_2)$. By Continuity 3, there exists $\overline{z}_i \in (x_i, \frac{x_i+x_{i+1}}{2})$ with $(x_i, p_{i,2}) \succ (\overline{z}_i, \underline{c}_2)$. Define $z_{i,j} = z_{i,1} + \frac{j-1}{t_i}(\overline{z}_i - z_{i,1})$ for all $j = 1, ..., t_i$. Clearly, $z_{i,1} = x_i < z_{i,2} < \cdots < z_{i,t_i} < \overline{z}_i$. By Monotonicity and Lemma A.2.1, I can find $\hat{z}_{i,j}$ for $j \geq 1$ with $(z_{i,j}, \hat{z}_{i,j}) \sim (x_i, p_{i,j})$. Then I define $d^i \in \mathcal{D}^*$ such that for any $x \notin \{z_{i,j}\}_{j=1}^{t_i}$, $q \in \mathcal{L}^0(X_2)$, $d^i(x, q) = d^{i-1}(x, q)$, and for each $j = 1, ..., t_i$, $d^1(z_{i,j}, \hat{z}_{i,j}) = d(x_i, p_{i,j})$ and $d^i(z_{i,j}, q) = 0$ for $q \neq \hat{z}_{i,j}$.

The algorithm terminates at $i = N < +\infty$ and $\text{supp}(d_1^N) = \bigcup_{k=1}^N \{z_{k,1}, \cdots, z_{k,t_k}\}$. For each $z_{k,i} \in \text{supp}(d_1^N)$ with $1 \leq i \leq t_k$ and $1 \leq k \leq N$, there exists $\hat{z}_{k,i}$ with $(z_{k,i}, \hat{z}_{k,i}) \sim (x_k, p_{k,i})$ and $d^N(z_{k,i}, \hat{z}_{k,i}) = d(x_k, p_{k,i})$. Moreover, $\{z_{k,i}\}$ admits a lexicographic order, that is, $z_{k,i} < z_{k',i'}$ if $k < k'$ or $k = k'$ and $i < i'$. This implies $d^N \in \hat{\mathcal{D}}^*$. In this way, I have constructed a mapping $h : \mathcal{D}^* \backslash \hat{\mathcal{D}}^* \to \hat{\mathcal{D}}^*$ where $h(d) = d^N$. Ac-

tually, the above algorithm also applies for $d \in \hat{\mathcal{D}}^*$ and hence I can define h on $\mathcal{D}^*$. Clearly $h(h(d)) = h(d)$ for all $d \in \mathcal{D}^*$. Then I define $\succsim^*$ on $\mathcal{D}^*$ such that $\succsim^*$ agrees with $\succsim'$ on $\hat{\mathcal{D}}^*$ and $d \sim^* h(d)$ for $d \in \mathcal{D}^* \backslash \hat{\mathcal{D}}^*$, that is, for each $d, d' \in \mathcal{D}^*$, $d \succsim^* d'$ if and only if $h(d) \succsim' h(d')$.

To verify that $\succsim^*$ is well-defined, we need to argue that the selection of $\{z_{k,i}\}$ does not affect the definition of $\succsim^*$. Consider two constructions h and $\hat{h}$. For each $d = \sum_{k,i} d(x_k, p_{k,i})\delta_{(x_k,p_{k,i})}$, $h(d) = \sum_{k,i} d(x_k, p_{k,i})\delta_{(z_{k,i},\hat{z}_{k,i})}$ and $\hat{h}(d) = \sum_{k,i} d(x_k, p_{k,i})\delta_{(z'_{k,i},\hat{z}'_{k,i})}$ such that $z_{k,i} \neq z_{k',i'}$, $z'_{k,i} \neq z'_{k',i'}$ for all $(k,i) \neq (k',i')$ and for all k, i, $(z'_{k,i}, \hat{z}'_{k,i}) \sim (z_{k,i}, \hat{z}_{k,i}) \sim (x_k, p_{k,i})$. By Corollary A.3.11 and the definition of $\succsim'$, $h(d) \sim' \hat{h}(d)$. Hence the definition of $\succsim^*$ is not affected by the construction of h. As a direct corollary, $h(d) \sim' d$ for each $d \in \hat{\mathcal{D}}^*$, although possibly $h(d) \neq d$.

Step 3. For $d, d' \in \mathcal{D}^* = \mathcal{L}^0(X_1 \times \mathcal{L}^0(X_2))$ and $\alpha \in (0,1)$, I define the α-mixture of d and d' as

$$[\alpha d + (1-\alpha)d'](x,q) = \alpha d(x,q) + (1-\alpha)d'(x,q), \forall\ \alpha \in (0,1), (x,q) \in X_1 \times \mathcal{L}^0(X_2).$$

Then $\mathcal{D}^*$ is a mixture space. I want to show that $\succsim^*$ satisfies the vNM independence axiom and mixture continuity on $\mathcal{D}^*$. I start with an intermediate result for the vNM independence axiom by extending Corollary A.3.11 to $\mathcal{D}^*$.

Lemma A.3.12 *Suppose $\lambda_i > 0$ for all i and $\sum_{i=1}^n \lambda_i = 1$. For $d^1 = \sum_{i=1}^n \lambda_i \delta_{(x_i,p_i)}$, $d^2 = \sum_{i=1}^n \lambda_i \delta_{(y_i,q_i)}$ with $(x_i,p_i) \sim (y_i,q_i)$ for all i, then $d^1 \sim^* d^2$.*

Proof of Lemma A.3.12. I first assume that $(x_i,p_i) \neq (x_j,p_j)$ and $(y_i,q_i) \neq (y_j,q_j)$ for all $i \neq j$. By definition of h, $h(d^1) = \sum_{i=1}^n \lambda_i \delta_{(z_i^1,\hat{z}_i^1)}$ with $z_i^1 \neq z_j^1$ for $i \neq j$ and $(z_i^1, \hat{z}_i^1) \sim (x_i, p_i)$ for each $i = 1, ..., n$. Similarly, $h(d^2) = \sum_{i=1}^n \lambda_i \delta_{(z_i^2,\hat{z}_i^2)}$ with $z_i^2 \neq z_j^2$ for $i \neq j$ and $(z_i^2, \hat{z}_i^2) \sim (x_i, p_i)$ for each $i = 1, ..., n$. Then $h(d^1) \sim^* d^1$ and $h(d^2) \sim^* d^2$. By Corollary A.3.11, $h(d^1) \sim^* h(d^2)$ and hence $d^1 \sim^* d^2$.

Now I allow $(x_i,p_i) = (x_j,p_j)$ or $(y_i,q_i) = (y_j,q_j)$ with some $i \neq j$. If $n = 1$, then the result is trivial. Suppose $n \geq 2$. Relabel the subscripts so that $(x_i,p_i) = (\underline{c}_1, \underline{c}_2)$ for $i \leq k$ for some $k \geq 0$ and $(x_i,p_i) \succ (\underline{c}_1, \underline{c}_2)$ for $i > k$. Since $(x_i,p_i) \sim (y_i,q_i)$, I

know that $(y_i, q_i) = (\underline{c}_1, \underline{c}_2)$ for $i \leq k$ and $(y_i, q_i) \succ (\underline{c}_1, \underline{c}_2)$ for $i > k$. Then

$$d^1 = [\sum_{i=1}^{k} \lambda_i]\delta_{(\underline{c}_1, \underline{c}_2)} + \sum_{i=k+1}^{n} \lambda_i \delta_{(x_i, p_i)} \ , \ d^2 = [\sum_{i=1}^{k} \lambda_i]\delta_{(\underline{c}_1, \underline{c}_2)} + \sum_{i=k+1}^{n} \lambda_i \delta_{(y_i, q_i)}.$$

This implies that I can assume $(x_i, p_i) \neq (\underline{c}_1, \underline{c}_2)$ for all $i \geq 2$. By a similar argument, I can assume $(x_i, p_i) \neq (\overline{c}_1, \overline{c}_2)$ for all $i < n$.

Without loss of generality, I can further assume $\overline{c}_1 > x_i > \underline{c}_1$ for all $2 \leq i \leq n-1$ since I can construct $\hat{d}^1 \sim d^1$ by replacing $(\underline{c}_1, p_i) \neq (\underline{c}_1, \underline{c}_2)$ with (a, q) and $(\overline{c}_1, p_i) \neq (\overline{c}_1, \overline{c}_2)$ with (b, q') for some $a, b \in (\underline{c}_1, \overline{c}_1)$.

Suppose that $(x_1, p_1) \succ (\underline{c}_1, \underline{c}_2)$ and $(x_n, p_n) \prec (\overline{c}_1, \overline{c}_2)$. By relabeling, there exists a partition of $\{1, ..., n\}$ as $\Big\{\{1, ..., t_1\}, \cdots, \{t_{k-1}+1, ..., n\}\Big\}$ such that $(x_i, p_i) = (x_j, p_j)$ for all $t_l + 1 \leq i, j \leq t_{l+1}$ with $0 \leq l \leq k-1$ and $t_0 = 0, t_k = n$. For $l = 0$, that is, $1 \leq i \leq t_1$, by Continuity and Lemma A.2.1, I can find $z_i > \underline{c}_1$, $\hat{z}_i \geq \underline{c}_2$ with $(z_i, \hat{z}_i) \sim (x_i, p_i) = (x_1, p_1)$ for all $i = 1, ..., t_1$ and $z_i \neq z_j$ for all $i \neq j$. By applying Lemma A.3.9 repeatedly,

$$\sum_{i=1}^{t_1} \frac{\lambda_i}{\sum_{j=1}^{t_1} \lambda_j}(z_i, \hat{z}_i) \sim (x_1, p_1).$$

The same result holds for $l = 1, ..., k-1$.

Recall that $d^1 = \sum_{i=1}^{n} \lambda_i \delta_{(x_i, p_i)} = \sum_{l=0}^{k-1} (\sum_{i=t_l+1}^{t_{l+1}} \lambda_i)\delta_{(x_i, p_i)}$. By definition of h, I can find $h(d^1) \sim^* d^1$ with $h(d^1) \in \hat{\mathcal{D}}^*$. Denote $h(d^1) = \sum_{l=0}^{k-1} \hat{\lambda}_{l+1} \delta_{(x'_{l+1}, \hat{x}'_{l+1})}$, where $\hat{\lambda}_{l+1} = \sum_{i=t_l+1}^{t_{l+1}} \lambda_i$ and $(x'_{l+1}, \hat{x}'_{l+1}) \sim (x_{t_l+1}, p_{t_l+1}) \sim \sum_{i=t_l+1}^{t_{l+1}} \frac{\lambda_i}{\sum_{j=t_l+1}^{t_{l+1}} \lambda_j}(z_i, \hat{z}_i)$ for each l. Denote $R_{l+1} = \sum_{i=t_l+1}^{t_{l+1}} \frac{\lambda_i}{\sum_{j=t_l+1}^{t_{l+1}} \lambda_j}(z_i, \hat{z}_i)$. Note that R_l and $R_{l'}$ are compatible, $(x'_l, \hat{x}'_l)$ and $(x'_{l'}, \hat{x}'_{l'})$ are compatible for all $l \neq l'$. By Lemma A.3.9 and the definition of $\succsim'$,

$$d^1 \sim^* h(d^1) = \sum_{l=1}^{k} \hat{\lambda}_l \delta_{(x'_l, \hat{x}'_l)} \sim' \sum_{l=1}^{k} h(\hat{\lambda}_l R_l) = \sum_{i=1}^{n} \lambda_i \delta_{(z_i, \hat{z}_i)} \in \hat{\mathcal{D}}^*.$$

Similarly, we can find $z'_i, \hat{z}'_i$ for $i = 1, ..., n$ such that $z'_i \neq z'_j$ for all $i \neq j$, $(z'_i, \hat{z}'_i) \sim (y_i, q_i)$ for each i and

$$d^2 \sim^* \sum_{i=1}^{n} \lambda_i \delta_{(z'_i, \hat{z}'_i)} \in \hat{\mathcal{D}}^*.$$

By Corollary A.3.11,

$$d^2 \sim^* \sum_{i=1}^{n} \lambda_i \delta_{(z_i', \hat{z}_i')} \sim^* \sum_{i=1}^{n} \lambda_i \delta_{(z_i, \hat{z}_i)} \sim^* d^1.$$

Now we turn to the case where $(x_1, p_1) = (\underline{c}_1, \underline{c}_2)$ or $(x_n, p_n) = (\overline{c}_1, \overline{c}_2)$ or both. This implies $(y_1, q_1) = (\underline{c}_1, \underline{c}_2)$ or $(y_n, q_n) = (\overline{c}_1, \overline{c}_2)$ or both. If $\lambda_1 + \lambda_n = 1$, then the result is trivial as $n = 2$ and I am back to the special case where $(x_i, p_i) \neq (x_j, p_j), (y_i, q_i) \neq (y_j, q_j)$ for all $i \neq j$. By Monotonicity, $(\overline{c}_1, \overline{c}_2) \succ (x_i, p_i) \succ (\underline{c}_1, \underline{c}_2), (\overline{c}_1, \overline{c}_2) \succ (y_i, q_i) \succ (\underline{c}_1, \underline{c}_2)$ for all $2 \leq i \leq n-2$. Define

$$\hat{d}^1 = \frac{1}{1 - \lambda_1 - \lambda_n} \sum_{i=2}^{n} \lambda_i \delta_{(x_i, p_i)} \ , \ \hat{d}^2 = \frac{1}{1 - \lambda_1 - \lambda_n} \sum_{i=2}^{n} \lambda_i \delta_{(y_i, q_i)}.$$

Then we return to the previous case and $\hat{d}^1 \sim^* \hat{d}^2$. By Lemma A.3.9 and the definition of $\succsim'$, $h(d^1) = \lambda_1 \delta_{(0,0)} + \lambda_n \delta_{(\overline{c}_1, \overline{c}_2)} + (1 - \lambda_1 - \lambda_n) h(\hat{d}^1) \sim' \lambda_1 \delta_{(0,0)} + \lambda_n \delta_{(\overline{c}_1, \overline{c}_2)} + (1 - \lambda_1 - \lambda_n) h(\hat{d}^2) = h(d^2)$. Thus, by definition of $\succsim^*$, $d^1 \sim^* d^2$. ■

Now I check the conditions of the Mixture Space Theorem for $\succsim^*$. First, I claim that $\succsim^*$ satisfies the vNM independence axiom, that is, for $d, d', d'' \in \mathcal{D}^*$ and $\alpha \in (0, 1)$,

$$d \succ^* d' \Longrightarrow \alpha d + (1 - \alpha) d'' \succ^* \alpha d' + (1 - \alpha) d''.$$

To see this, fix $d, d', d'' \in \mathcal{D}^*$ with $d \succ^* d'$ and $\alpha \in (0, 1)$. Denote $d = \sum_{i=1}^{n} \lambda_i \delta_{(x_i, p_i)}$ and $d'' = \sum_{j=1}^{m} \eta_j \delta_{(y_j, q_j)}$. As $\text{supp}(d_1) \cup \text{supp}(d_1'')$ is finite, for any i such that $(x_i, p_i) \neq (\underline{c}_1, \underline{c}_2)$ and $(\overline{c}_1, \overline{c}_2)$, I can find $(z_i^d, \hat{z}_i^d) \sim (x_i, p_i)$; for any j such that $(y_j, q_j) \neq (\underline{c}_1, \underline{c}_2)$ and $(\overline{c}_1, \overline{c}_2)$, we can find $(z_j^{d''}, \hat{z}_j^{d''}) \sim (y_j, q_j)$. Moreover, we can require that $\overline{c}_1 > z_i^d, z_j^{d''} > \underline{c}_1$, $z_i^d \neq z_j^d, z_i^{d''} \neq z_j^{d''}$ for all $i \neq j$ and $z_i^d \neq z_j^{d''}$ for all i, j. For any i, j with $(x_i, p_i) = (\underline{c}_1, \underline{c}_2)$ or $(y_j, q_j) = (\underline{c}_1, \underline{c}_2)$, denote $z_i^d = z_j^{d''} = \underline{c}_1$ and $\hat{z}_i^d = \hat{z}_j^{d''} = \underline{c}_2$. For any i, j with $(x_i, p_i) = (\overline{c}_1, \overline{c}_2)$ or $(y_j, q_j) = (\overline{c}_1, \overline{c}_2)$, denote $z_i^d = z_j^{d''} = \overline{c}_1$ and $\hat{z}_i^d = \hat{z}_j^{d''} = \overline{c}_2$. Then Lemma A.3.12 assures that

$$\alpha d + (1 - \alpha) d'' \sim^* \alpha \hat{d} + (1 - \alpha) \hat{d}'',$$

where $\hat{d} = \sum_{i=1}^{n} \lambda_i \delta_{(z_i^d, \hat{z}_i^d)} \sim^* d$ and $\hat{d}'' = \sum_{j=1}^{m} \eta_j \delta_{(z_i^{d''}, \hat{z}_i^{d''})} \sim^* d''$. As $\hat{d}, \hat{d}'' \in \hat{\mathcal{D}}^*$, I can denote $P = f^{-1}[\hat{d}]$ and $R = f^{-1}[\hat{d}'']$. Easy to see that P and R are weakly compatible, and $\alpha d + (1-\alpha)d'' \sim^* \alpha f(P) + (1-\alpha) f(R)$ for any $\alpha \in (0,1)$.

Similarly, we can construct Q and S such that $d' \sim^* f(Q)$, $d'' \sim^* f(S)$, Q and S are weakly compatible and $\alpha d' + (1-\alpha)d'' \sim^* \alpha f(Q) + (1-\alpha) f(S)$ for any $\alpha \in (0,1)$. Note that $d \succ^* d'$ if and only if $P \succ Q$ and $d'' = d''$ implies $R \sim S$. By Lemma A.3.9, $\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S$ for any $\alpha \in (0,1)$. It is easy to verify that $f(\alpha P + (1-\alpha R)) = \alpha f(P) + (1-\alpha) f(R)$ and $f(\alpha Q + (1-\alpha S)) = \alpha f(Q) + (1-\alpha) f(S)$ since P and R are weakly compatible and Q and S are weakly compatible. Thus for each $\alpha \in (0,1)$,

$$
\begin{aligned}
&\alpha P + (1-\alpha)R \succ \alpha Q + (1-\alpha)S \\
\Longleftrightarrow &f(\alpha P + (1-\alpha)R) \succ^* f(\alpha Q + (1-\alpha)S) \\
\Longleftrightarrow &\alpha f(P) + (1-\alpha) f(R) \succ^* \alpha f(Q) + (1-\alpha) f(S) \\
\Longleftrightarrow &\alpha d + (1-\alpha)d'' \succ^* \alpha d' + (1-\alpha)d''.
\end{aligned}
$$

Hence $\succsim^*$ satisfies the vNM independence axiom on $\mathcal{D}^*$.

Then I show the mixture continuity of $\succsim^*$ on $\mathcal{D}^*$. For any $d, d', d'' \in \mathcal{D}^*$, using the above arguments, I can find $P, Q, R \in \mathcal{P}$ such that $f(P) \sim^* d$, $f(Q) \sim^* d'$, $f(R) \sim^* d''$ and for each $\alpha \in (0,1)$, $\alpha P + (1-\alpha)Q \in \mathcal{P}$, $f(\alpha P + (1-\alpha)Q) = \alpha f(P) + (1-\alpha) f(Q)$ and $\alpha d + (1-\alpha)d' \sim^* \alpha f(P) + (1-\alpha) f(Q)$. Then

$$
\begin{aligned}
A &= \left\{\alpha \in [0,1] : \alpha d + (1-\alpha)d' \succ^* d''\right\} \\
&= \left\{\alpha \in [0,1] : \alpha f(P) + (1-\alpha) f(Q) \succ^* f(R)\right\} \\
&= \left\{\alpha \in [0,1] : f(\alpha P + (1-\alpha)Q) \succ^* f(R)\right\} \\
&= \left\{\alpha \in [0,1] : \alpha P + (1-\alpha)Q \succ R\right\}.
\end{aligned}
$$

By mixture continuity of $\succsim$ on $\mathcal{P}$, I know A is open in $[0,1]$. Similarly, $\left\{\alpha \in [0,1] : \alpha d + (1-\alpha)d' \prec^* d''\right\}$ is open in $[0,1]$. Thus, $\succsim^*$ satisfies mixture continuity on $\mathcal{D}^*$.

By the Mixture Space Theorem, $\succsim^*$ on $\mathcal{D}^*$ admits an EU representation U with a utility index $w_D : X_1 \times \mathcal{L}^0(X_2) \to \mathbb{R}$. That is, for each $d \in \mathcal{D}^*$

$$U(d) = \sum_{x,p} w_D(x,p)d(x,p).$$

Moreover, w_D is unique up to a positive affine transformation.

Step 4. Since $\succsim^*$ extends $\succsim'$ from $\hat{\mathcal{D}}^*$ to $\mathcal{D}^*$, $d \succsim' d'$ if and only if $f^{-1}(d) \succsim f^{-1}(d')$ for all $d, d' \in \hat{\mathcal{D}}^*$ and f is bijective, $\succsim$ on $\mathcal{P}$ is represented by

$$V(P) = \sum_x P_1(x) w_D(x, P_{2|x}), \forall\ P \in \mathcal{P}.$$

Then I want to show that the above utility function can be rewritten as (A.6).

First, by Lemma A.2.1, for each $x \in X_1$, $\succsim_{2|x}$ on $\{\delta_x\} \times \mathcal{L}^0(X_2)$ admits an EU representation with some continuous and strictly monotonic utility index v_x. Then there exists a continuous and strictly monotonic function ϕ_x such that for all $p \in \mathcal{L}^0(X_2)$,

$$V(x,p) = w_D(x,p) = \phi_x(CE_{v_x}(p)).$$

Define $\hat{u}_D : X_1 \times X_2 \to \mathbb{R}$ as $\hat{w}(x,y) = \phi_x(y) = w_D(x,y)$ for all $(x,y) \in X_1 \times X_2$. Then the utility function can be rewritten as

$$V(P) = \sum_x \hat{u}_D(x, CE_{v_x}(P_{2|x}))P_1(x), \forall\ P \in \mathcal{P}$$

where v_x is continuous and strictly monotonic for each $x \in X_1$. This is exactly (A.6). By Monotonicity, $\hat{u}_D$ is strictly monotonic. $\hat{u}_D$ is bounded since $\hat{u}_D(\underline{c}_1, \underline{c}_2) \leq \hat{u}_D(x,y) \leq \hat{u}_D(\overline{c}_1, \overline{c}_2)$ for all $(x,y) \in X_1 \times X_2$. Moreover, $\hat{u}_D$ is unique up to a positive affine transformation.

The final step is to verify that $\hat{u}_D$ is continuous. Suppose by contradiction that $\hat{u}_D$ is not continuous, then I can find $(x,y) \in X_1 \times X_2$ and a sequence $(x_n, y_n) \to (x,y)$ such that $\lim_{n\to\infty} \hat{u}_D(x_n, y_n) \neq \hat{u}_D(x,y)$. As $\hat{u}_D$ is bounded, I can find a subsequence of $\{(x_n, y_n)\}_{n\geq 1}$ (still denoted as $\{(x_n, y_n)\}_{n\geq 1}$) such that $\lim_{n\to\infty} \hat{u}_D(x_n, y_n) = a \neq$

$b = \hat{u}_D(x, y)$. Without loss of generality, assume that $a < b$. Note that for $x_1 \neq x_2$ and $\eta \in (0, 1)$, $V(\eta(x_1, y_1) + (1 - \eta)(x_2, y_2)) = \eta V(x_1, y_1) + (1 - \eta)V(x_2, y_2)$. I can find some $P \in \mathcal{P}$ with $V(P) \in (a, b)$. As $\lim_{n\to\infty} \hat{u}_D(x_n, y_n) = a < V(P)$, for n large enough, I have $\hat{u}_D(x_n, y_n) < V(P) < b$, that is, $(x_n, y_n) \prec P \prec (x, y)$. When n goes to infinity, $(x_n, y_n) \xrightarrow{w} (x, y)$. However, $(x, y) \precsim P$ by Continuity, a contradiction. Hence $\hat{u}_D$ is continuous. This completd the proof of Lemma A.3.10. ■

Recall that $\hat{\succsim}$ satisfies Correlation Neutrality and agrees with $\succsim$ on $\hat{\mathcal{P}}$. Corollary A.3.1 states that $\hat{\succsim}$ can be represented by one of the following types of utility functions: V^{CN}, V^{ANB-CN} or V^{FNB}. If $\succsim$ satisfies Correlation Neutrality, then $\succsim = \hat{\succsim}$ and it is also represented by one of V^{CN}, V^{ANB-CN} or V^{FNB}. If instead $\succsim$ violates Correlation Neutrality, then Lemma A.3.10 states that $\succsim$ is represented by the utility function (A.6)

$$U(P) = \sum_x \hat{u}_D(x, CE_{v_x}(P_{2|x}))P_1(x),$$

where $\hat{u}_D$ and v_x for each $x \in X_1$ are continuous, strictly monotonic and unique up to a positive affine transformation. Hence I can exploit the consistency of $\succsim$ and $\hat{\succsim}$ on $\hat{\mathcal{P}}$ to refine the utility function (A.6). Note that on $\hat{\mathcal{P}}$, (A.6) becomes

$$U(p, q) = \sum_x \hat{u}_D(x, CE_{v_x}(q))p(x),$$

Case 1: Suppose $\hat{\succsim}$ is represented by V^{CN}. Then for each $(p, q) \in \hat{\mathcal{P}}$,

$$V^{CN}(p, q) = \sum_{x,y} w(x, y)p(x)q(y)$$

where w is continuous, strictly monotonic and unique up to a positive affine transformation. Then it must be the case that u_D is a positive affine transformation of w and v_x is a positive affine transformation of $w(x, \cdot)$ for all $x \in X_2$. Thus $\succsim$ on $\mathcal{P}$ is represented by

$$V^{EU}(P) = \sum_x w(x, CE_{w(x,\cdot)}(P_{2|x}))P_1(x) = \sum_{x,y} w(x, y)P_1(x)P_{2|x}(y) = \sum_{x,y} w(x, y)P(x, y).$$

Case 2: Suppose that $\hat{\succsim}$ is represented by V^{BIB-CN} or V^{FNB}. Then for each $x \in X_1$, $\succsim_{2|x}$ is admits an EU representation with some index v_2, which is independent of x. This suggests that v_x must be a positive affine transformation of v_2 for all $x \in X_1$. Thus $\succsim$ on $\mathcal{P}$ is represented by

$$V^{ANB}(P) = \sum_x \hat{u}_D(x, CE_v(P_{2|x}))P_1(x).$$

To summarize, if $\succsim$ satisfies Weak Order, Monotonicity, Continuity, Independence* and Correlation Consistency, then $\succsim$ an be represented by one of the following types of utility functions: V^{EU}, V^{ANB}, V^{CN} , V^{ANB-CN} or V^{FNB}.

A.4 Proof of Proposition 1.4.3

The arguments for uniqueness of utility indices in the representations are already contained in the proofs of Theorem 1.4.1 and Theorem 1.4.2.

A.5 Proof of Corollary 1.6.2

For necessity, Symmetry is guaranteed since both $\succsim_1$ and $\succsim_2$ admit an expected utility representation with utility index u. Stochastic Dominance without Risk is satisfied since for each $x, y, x', y' \in X$, $(x, y) \succsim (x', y')$ if and only if $u(x+y) \geq u(x'+y')$ if and only if $x + y \geq x' + y'$. The last equivalence holds as u is strictly monotonic.

Now I check sufficiency. For each utility function, Stochastic Dominance without Risk implies that $w(x, y)$ can be rewritten as $\phi(x + y)$ for each $(x, y) \in Z$, where ϕ is strictly monotonic and continuous. By Symmetry, the ϕ agrees with u in utility functions V^{ANB} and V^{ANB-CN} and v agrees with u in V^{FNB}. Since w is unique up to a monotonic transformation in V^{FNB}, the five utility functions can be written as in Corollary 1.6.2.

A.6 Proof of Proposition 1.6.3

I first prove (v). For any u continuous and strictly monotonic and $P \in \mathcal{P}$, $V^{EU}(P) = \sum_{x,y} u(x+y)P(x,y) = \sum_z u(z)f[P](z)$. By monotonicity of u, $V^{EU}(P) > V^{EU}(Q)$ if $f[P] \succ_{FOSD} f[Q]$. Hence EU satisfies stochastic dominance on $\mathcal{C}_4 = \mathcal{P}$.

(i). By Corollary 1.6.2, FNB satisfies Stochastic Dominance without Risk and hence it satisfies stochastic dominance on $\mathcal{C}_1$. Now I provide an example that FNB might violate stochastic dominance on $\mathcal{C}_2$. Consider $P = (\frac{1}{2}\delta_{101} + \frac{1}{2}\delta_0, \delta_{100})$ and $Q = (\frac{1}{2}\delta_{200} + \frac{1}{2}\delta_{100}, \delta_0)$. Clearly, $f[P] \succ_{FOSD} f[Q]$ and $P, Q \in \mathcal{C}_2$. Assume that $u(x) = \sqrt{x}$ for $x \in X \cap \mathbb{R}_+$ and $u(x) = -\sqrt{-x}$ for $x \in X \cap \mathbb{R}_-$. Then $V^{FNB}(P) = (101/4 + 100)^{1/2} < V^{FNB}(Q) = 5(1 + \sqrt{2})$ and FNB does not satisfy stochastic dominance on $\mathcal{C}_2$.

(ii). For any u continuous and strictly monotonic and $(p, y) \in \mathcal{C}_2$,

$$V^{ANB-CN}(p, y) = \sum_{x \in X} u(x+y)p(x) = V^{EU}(p, y).$$

Then (v) implies that ANB-CN satisfies stochastic dominance on $\mathcal{C}_2$.

To show that ANB-CN might violate stochastic dominance on $\mathcal{C}_3$, consider $P = \frac{1}{2}(0, 100) + \frac{1}{2}(100, 0)$ and $Q = (99, 0)$. Clearly, $f[P] \succ_{FOSD} f[Q]$ and $P, Q \in \mathcal{C}_3$. Assume that $u(x) = \sqrt{x}$ for $x \in X \cap \mathbb{R}_+$ and $u(x) = -\sqrt{-x}$ for $x \in X \cap \mathbb{R}_-$. Then $V^{ANB-CN}(P) = \frac{1}{2}(5 + 5\sqrt{5}) < V^{FNB}(Q) = \sqrt{99}$ and ANB-CN does not satisfy stochastic dominance on $\mathcal{C}_3$.

To show that ANB-CN might violate stochastic dominance on $\mathcal{C}_{3'}$, consider $P = (\delta_{100}, \frac{1}{2}\delta_{101} + \frac{1}{2}\delta_0)$ and $Q = (\delta_0, \frac{1}{2}\delta_{200} + \frac{1}{2}\delta_{100})$. Clearly, $f[P] \succ_{FOSD} f[Q]$ and $P, Q \in \mathcal{C}_{3'}$. Assume that $u(x) = \sqrt{x}$ for $x \in X \cap \mathbb{R}_+$ and $u(x) = -\sqrt{-x}$ for $x \in X \cap \mathbb{R}_-$. Then $V^{FNB}(P) = (101/4 + 100)^{1/2} < V^{FNB}(Q) = 5(1 + \sqrt{2})$ and ANB-CN does not satisfy stochastic dominance on $\mathcal{C}_{3'}$.

(iii). For any u continuous and strictly monotonic and $P \in \mathcal{C}_3$, denote $P_{2|x} = z_x$ for each $P_1(x) > 0$, and I have

$$V^{ANB}(P) = \sum_{x \in X} u(x + z_x)p(x) = V^{EU}(P).$$

Then (v) implies that ANB satisfies stochastic dominance on $\mathcal{C}_3$. The example that ANB does not satisfy stochastic dominance on $\mathcal{C}_{3'}$ is the same as that of ANB-CN.

(iv). For any u continuous and strictly monotonic and $(p, q) \in \mathcal{C}_{3'}$,

$$V^{CN}(p, q) = \sum_{x,y \in X} u(x + y)p(x)q(y) = V^{EU}(p, q).$$

Then (v) implies that CN satisfies stochastic dominance on $\mathcal{C}_{3'}$. The example that CN does not satisfy stochastic dominance on $\mathcal{C}_3$ is the same as that of ANB-CN.

A.7 Proof of Corollary 1.7.1

The proof is similar to the proof of Corollary 1.6.2 and omitted.

A.8 Proof of Lemma 1.7.5

s^{EU} satisfies the following first order condition:

$$\frac{1}{y_1 - s} = \beta R \frac{\sum_i \pi_i (y_{2,i} + Rs)^{\beta\alpha - 1}}{\sum_i \pi_i (y_{2,i} + Rs)^{\beta\alpha}}. \tag{A.7}$$

Similarly, s^{NB} satisfies the following first order condition:

$$\frac{1}{y_1 - s} = \beta R \frac{\sum_i \pi_i (y_{2,i} + Rs)^{\alpha - 1}}{\sum_i \pi_i (y_{2,i} + Rs)^{\alpha}}. \tag{A.8}$$

For any $\alpha < 1$ and $\alpha \neq 0$, one can easily show that expressions on the right hand side (RHS) of equations (A.7) and (A.8) are strictly decreasing in s, while the expressions on the left hand side (LHS) are strictly increasing in s. This implies that

both s^{EU} and s^{NB} are unique and $\frac{y_{2,i}}{R} < s^{NB}, s^{EU} < y_1$ for all i. Moreover, if

$$\beta R \frac{\sum_i \pi_i (y_{2,i} + Rs)^{\beta\alpha - 1}}{\sum_i \pi_i (y_{2,i} + Rs)^{\beta\alpha}} < \beta R \frac{\sum_i \pi_i (y_{2,i} + Rs)^{\alpha - 1}}{\sum_i \pi_i (y_{2,i} + Rs)^{\alpha}} \tag{A.9}$$

for any interior s, i.e., $\frac{y_{2,i}}{R} < s < y_1$ for all i, then I claim that $s^{NB} > s^{EU}$. To see this, suppose by contradiction that $s^{NB} \leq s^{EU}$. Then

$$\begin{aligned}
\frac{1}{y_1 - s^{EU}} &= \beta R \frac{\sum_i \pi_i (y_{2,i} + Rs^{EU})^{\beta\alpha - 1}}{\sum_i \pi_i (y_{2,i} + Rs^{EU})^{\beta\alpha}} \\
&\leq \beta R \frac{\sum_i \pi_i (y_{2,i} + Rs^{NB})^{\beta\alpha - 1}}{\sum_i \pi_i (y_{2,i} + Rs^{NB})^{\beta\alpha}} \\
&< \beta R \frac{\sum_i \pi_i (y_{2,i} + Rs^{NB})^{\alpha - 1}}{\sum_i \pi_i (y_{2,i} + Rs^{NB})^{\alpha}} \\
&= \frac{1}{y_1 - s^{NB}}.
\end{aligned}$$

The first and the last inequalities come from (A.7) and (A.8). The second inequality holds since $s^{NB} \leq s^{EU}$ and the RHS of (A.7) is decreasing in s. The third inequality follows from (A.9). Hence, I have $s^{EU} < s^{NB}$, which leads to a contradiction.

It remains to be shown that (A.9) holds for all $\frac{y_{2,i}}{R} < s < y_1$ for all i. Note that (A.9) is equivalent to

$$\sum_{i,j} \pi_i \pi_j (y_{2,i} + Rs)^{\beta\alpha - 1} (y_{2,j} + Rs)^{\alpha} < \sum_{i,j} \pi_i \pi_j (y_{2,i} + Rs)^{\alpha - 1} (y_{2,j} + Rs)^{\beta\alpha}.$$

Both sides of the inequality consist of N^2 elements, each indexed by (i, j). For $i = j$, the elements on the two sides cancel out. Since $\pi_i > 0$ and $\frac{y_{2,i}}{R} < s < y_1$ for all $i = 1, ..., N$, it suffices to show that for each $i \neq j$,

$$\begin{aligned}
&(y_{2,i} + Rs)^{\beta\alpha - 1} (y_{2,j} + Rs)^{\alpha} + (y_{2,j} + Rs)^{\beta\alpha - 1} (y_{2,i} + Rs)^{\alpha} \\
&< (y_{2,i} + Rs)^{\alpha - 1} (y_{2,j} + Rs)^{\beta\alpha} + (y_{2,j} + Rs)^{\alpha - 1} (y_{2,i} + Rs)^{\beta\alpha} \\
\iff &(y_{2,j} - y_{2,i}) \Big((y_{2,j} + Rs)^{\alpha(1-\beta)} - (y_{2,i} + Rs)^{\alpha(1-\beta))} \Big) < 0.
\end{aligned}$$

Since $y_{2,j} \neq y_{2,i}$ if $i \neq j$ and $\beta \in (0,1)$, the above inequality holds if $\alpha < 0$. Hence, when $\alpha < 0$, $s^{NB} > s^{EU}$. Symmetrically, if $\alpha > 0$, then the inequality in (A.9) is reversed and thus $s^{NB} < s^{EU}$.

A.9 Proof of Lemma 1.7.6

Denote by s_i the saving level in state $i = l, h$. s_l^{EU} and s_h^{EU} satisfy the following first order condition:

$$\frac{y_{2,i} + Rs_i}{y_1 - s_i} = \beta R. \tag{A.10}$$

If the consumer's preference is represented by V^{CN}, then she chooses s_l and s_h to maximize

$$\max_{s_l, s_h} \frac{1}{4\alpha} \sum_{i,j=l.h} (y_1 - s_i)^{\alpha} (y_{2,j} + Rs_j)^{\beta\alpha}$$

s_l^{CN} and s_h^{CN} satisfy the following first order conditions:

$$(y_1 - s_l)^{\alpha-1}\Big[(y_{2,l} + Rs_l)^{\beta\alpha} + (y_{2,h} + Rs_h)^{\beta\alpha}\Big] = \beta R (y_{2,l} + Rs_l)^{\beta\alpha-1}\Big[(y_1 - s_l)^{\alpha} + (y_1 - s_h)^{\alpha}\Big], \tag{A.11}$$

$$(y_1 - s_h)^{\alpha-1}\Big[(y_{2,l} + Rs_l)^{\beta\alpha} + (y_{2,h} + Rs_h)^{\beta\alpha}\Big] = \beta R (y_{2,h} + Rs_h)^{\beta\alpha-1}\Big[(y_1 - s_l)^{\alpha} + (y_1 - s_h)^{\alpha}\Big]. \tag{A.12}$$

Divide equation (A.11) by (A.12) and I get

$$\frac{(y_1 - s_l)^{\alpha-1}}{(y_1 - s_h)^{\alpha-1}} = \frac{(y_{2,l} + Rs_l)^{\beta\alpha-1}}{(y_{2,h} + Rs_h)^{\beta\alpha-1}}. \tag{A.13}$$

First, I claim that $s_l^{CN} > s_h^{CN}$. Suppose by contradiction that $s_l^{CN} \leq s_h^{CN}$. Since $\alpha < 1, \beta\alpha < 1$ and $y_{2,l} < y_{2,h}$, easy to see that $y_1 - s_l^{CN} \geq y_1 - s_h^{CN}$ and $y_{2,l} + Rs_l^{CN} < y_{2,h} + Rs_h^{CN}$. Then

$$\frac{(y_1 - s_l^{CN})^{\alpha-1}}{(y_1 - s_h^{CN})^{\alpha-1}} \leq 1 < \frac{(y_{2,l} + Rs_l^{CN})^{\beta\alpha-1}}{(y_{2,h} + Rs_h^{CN})^{\beta\alpha-1}},$$

which contradicts with (A.13). Rewrite (A.11) and (A.12) as follows:

$$\frac{y_{2,l}+Rs_l^{CN}}{y_1-s_l^{CN}}=\beta R\frac{1+\frac{(y_1-s_h^{CN})^\alpha}{(y_1-s_l^{CN})^\alpha}}{1+\frac{(y_{2,h}+Rs_h^{CN})^{\beta\alpha}}{(y_{2,l}+Rs_l^{CN})^{\beta\alpha}}}, \tag{A.14}$$

$$\frac{y_{2,h}+Rs_h^{CN}}{y_1-s_h^{CN}}=\beta R\frac{1+\frac{(y_1-s_l^{CN})^\alpha}{(y_1-s_h^{CN})^\alpha}}{1+\frac{(y_{2,l}+Rs_l^{CN})^{\beta\alpha}}{(y_{2,h}+Rs_h^{CN})^{\beta\alpha}}}. \tag{A.15}$$

I focus on the comparison of s_l^{CN} and s_l^{EU}. The comparison of s_h^{CN} and s_h^{EU} is similar. The expressions on the LHS of (A.14) and (A.10) are strictly increasing in s_l. Like the proof of Lemma 1.7.5, I know that $s_l^{CN}>s_l^{EU}$ if and only if

$$\frac{(y_1-s_h^{CN})^\alpha}{(y_1-s_l^{CN})^\alpha}>\frac{(y_{2,h}+Rs_h^{CN})^{\beta\alpha}}{(y_{2,l}+Rs_l^{CN})^{\beta\alpha}}.$$

Since s_l^{CN} and s_h^{CN} satisfy (A.13), the above inequality reduces to

$$\frac{(y_{2,l}+Rs_l^{CN})}{(y_1-s_l^{CN})}>\frac{(y_{2,h}+Rs_h^{CN})}{(y_1-s_h^{CN})}. \tag{A.16}$$

Again by (A.13),

$$\Big(\frac{y_1-s_l^{CN}}{y_1-s_h^{CN}}\Big)^{\alpha(1-\beta)}=\Big(\frac{\frac{y_{2,h}+Rs_h^{CN}}{y_1-s_h^{CN}}}{\frac{y_{2,l}+Rs_l^{CN}}{y_1-s_l^{CN}}}\Big)^{1-\beta\alpha}.$$

Since $\alpha<1$, $\beta\in(0,1)$ and $s_l^{CN}>s_h^{CN}$, (A.16) holds if and only if $\alpha>0$. Hence, $s_l^{CN}>s_l^{EU}$ if $\alpha>0$ and $s_l^{CN}<s_l^{EU}$ if $\alpha<0$. Similarly, one can show by (A.15) that $s_h^{CN}<s_h^{EU}$ if $\alpha>0$ and $s_h^{CN}>s_h^{EU}$ if $\alpha<0$.

A.10 Proof of Proposition A.1.1

First, by Theorem 5.1 in Epstein and Zin (1989), we can find an optimal consumption plan $\{c_t^*\}_{t\geq1}$ for the EZ consumer with $J^{EZ}(W_1)=\hat{V}^{EZ}(\{c_t^*\}_{t\geq1})$. On the one hand, since $\rho>\alpha$, the EZ consumer has a preference for early resolution of risk. Then $\hat{V}^{ANB}(\{c_t\}_{t\geq1})\leq\hat{V}^{EZ}(\{c_t\}_{t\geq1})$ for each feasible consumption plan $\{c_t\}_{t\geq1}$ and $J^{ANB}(W_1)\leq J^{EZ}(W_1)$. Moreover, for a consumption plan $\{c_t\}_{t\geq1}$ where c_t is injec-

tive on Ω_t for each t, the consumption history contains the same information as state history. In this case, $\hat{V}^{ANB}(\{c_t\}_{t\geq 1}) = \hat{V}^{EZ}(\{c_t\}_{t\geq 1})$.

If c_t^* is injective for each $t \geq 2$, set $c_t^n \equiv c_t^*$ for each n and the results holds. Otherwise, I want to construct a sequence of consumption plans $\{c_t^n\}_{t\geq 1, n\geq 1}$ with injective consumption functions such that $c_1^n \to c_1^*, c_t^n(s^t) \to c_t^*(s^t)$ as n goes to infinity for each $t \geq 2$. By continuity of V^{EZ} and hence $\hat{V}^{EZ}$,

$$J^{ANB}(W_1) \geq \lim_{n\to\infty} \hat{V}^{ANB}(\{c_t^n\}_{t\geq 1}) = \hat{V}^{EZ}(\{c_t^*\}_{t\geq 1}) = J^{EZ}(W_1).$$

Hence $J^{ANB}(W_1) = J^{EZ}(W_1)$. It remains to construct the sequence of consumption plans. This is easy since $\bigcup_{t\geq 2} c_t^*(S^t)$ is countable (as the union of a countable number of countable sets is countable), while the space of consumption is a continuum.

A.11 Proof of Corollary A.1.2

The derivation of the Euler equation and the asset pricing formula for $\{c_t\}_{t\geq 1}$ can be found in Epstein and Zin (1991). The rest follows from the continuity of the two equations and $c_1^n \to c_1, c_t^n(s^t) \to c_t(s^t)$ as n goes to infinity for each $t \geq 2$.

Appendix B

Appendices for Chapter 2

B.1 Revelation Principle

We prove the revelation principle for the single-agent case in this section. The multi-agent case can be shown similarly. Let the message space M be a nonempty measurable space. We assume that the message space is the same for all single-mechanisms in order to guarantee that knowing the set of messages does not provide any information about which single-mechanism is used. Given M, we redefine single-mechanisms, multi-mechanisms and implementability of a choice correspondence. We then show that the revelation principle holds.

A single-mechanism is a tuple (G, T) where $G : M \to \Delta(X)$ and $T : M \to \mathbb{R}$. Under the single-mechanism (G, T), when the agent reports $m \in M$, the outcome is chosen according to the distribution $G(m)$ and the transfer paid by the agent is $T(m)$. A multi-mechanism is a nonempty set of single-mechanisms $(G_i, T_i)_{i \in I}$, where the MD commits to using some single-mechanism indexed by $i \in I$ but conceals which one to use. A reporting strategy of the agent is a map $\mu : \Theta \to \Delta(M)$.

Definition B.1.1 *A multi-mechanism $(G_i, T_i)_{i \in I}$ implements a choice correspondence F if there exists a reporting strategy μ such that the following holds.*

1. *(Optimality) For any $\theta \in \Theta$ and $\eta \in \Delta(M)$,*

$$\inf_{i\in I}\left\{\int_M\int_X u(\theta,x)G_i(m)[dx]\mu(\theta)[dm]-\int_M T_i(m)\mu(\theta)[dm]\right\}$$

$$\geq \inf_{i\in I}\left\{\int_M\int_X u(\theta,x)G_i(m)[dx]\eta[dm]-\int_M T_i(m)\eta[dm]\right\}.$$

2. *(Non-triviality) For any $\theta \in \Theta$,*

$$\inf_{i\in I}\left\{\int_M\int_X u(\theta,x)G_i(m)[dx]\mu(\theta)[dm]-\int_M T_i(m)\mu(\theta)[dm]\right\} > -\infty.$$

3. *(Consistency) For any $i \in I$ and $\theta \in \Theta$, $\int_M G_i(m)[F(\theta)]\mu(\theta)[dm] = 1$.*

A choice correspondence F is implementable if there is a multi-mechanism implementing it.

Now we show that any implementable choice correspondence F in terms of Definition B.1.1 is also implementable in terms of Definition 2.2.1. Fix the message space M, the multi-mechanism $(G_i, T_i)_{i\in I}$ and the reporting strategy μ. We construct the multi-mechanism $(g_i, t_i)_{i\in I}$ under the message space Θ as follows. For any $\theta \in \Theta$, $g_i(\theta)$ is a distribution over X such that for any measurable set $E \subseteq X$, $g_i(\theta)[E] = \int_M G_i(m)[E]\mu(\theta)[dm]$. For any $\theta \in \Theta$, $t_i(\theta) = \int_M T_i(m)\mu(\theta)[dm]$. One can check that the multi-mechanism $(g_i, t_i)_{i\in I}$ satisfies the conditions stated in Definition 2.2.1, and thus the revelation principle holds.

B.2 Proof of Theorem 2.3.2

Necessity. Suppose that F is implemented by $(g_i, t_i)_{i\in I}$. For any nonempty finite sequence of types $\{\theta_1, ..., \theta_n\} \subseteq \Theta$, truth-telling implies that for any $k \in \{1, ..., n\}$, we have

$$\begin{aligned}&\inf_{i\in I}\left\{\int_X u(\theta_k,x)g_i(\theta_k)[dx]-t_i(\theta_k)\right\}\\ \geq &\inf_{i\in I}\left\{\int_X u(\theta_k,x)g_i(\theta_{k+1})[dx]-t_i(\theta_{k+1})\right\}.\end{aligned} \tag{B.1}$$

i.e., a type θ_k agent cannot be better off by reporting θ_{k+1}. By non-triviality, the LHS of (B.1) is bounded below. Thus, for each $\epsilon > 0$, by the definition of infimum, there exists some $i(\epsilon, \theta_k, \theta_{k+1}) \in I$ such that

$$\begin{aligned} &\inf_{i \in I} \left\{ \int_X u(\theta_k, x) g_i(\theta_k)[dx] - t_i(\theta_k) \right\} \\ \geq &\int_X u(\theta_k, x) g_{i(\epsilon,\theta_k,\theta_{k+1})}(\theta_{k+1})[dx] - t_{i(\epsilon,\theta_k,\theta_{k+1})}(\theta_{k+1}) - \epsilon. \end{aligned} \tag{B.2}$$

Relaxing the LHS of (B.2) by replacing the infimum payoff by the payoff under the single mechanism $i(\theta_{k-1}, \theta_k, \epsilon)$, we have

$$\begin{aligned} &\int_X u(\theta_k, x) g_{i(\epsilon,\theta_{k-1},\theta_k)}(\theta_k)[dx] - t_{i(\epsilon,\theta_{k-1},\theta_k)}(\theta_k) \\ \geq &\int_X u(\theta_k, x) g_{i(\epsilon,\theta_k,\theta_{k+1})}(\theta_{k+1})[dx] - t_{i(\epsilon,\theta_k,\theta_{k+1})}(\theta_{k+1}) - \epsilon. \end{aligned} \tag{B.3}$$

By summing up the n inequalities in the form of (B.3) to eliminate the transfers, we get

$$n\epsilon + \sum_{k=1}^{n} \left\{ \int_X [u(\theta_{k+1}, x) - u(\theta_k, x)] g_{i(\epsilon,\theta_k,\theta_{k+1})}(\theta_{k+1})[dx] \right\} \geq 0. \tag{B.4}$$

By consistency, we know that

$$g_{i(\epsilon,\theta_k,\theta_{k+1})}(\theta_{k+1})[F(\theta_{k+1})] = 1.$$

Thus, inequality (B.4) implies that

$$n\epsilon + \sum_{k=1}^{n} \left\{ \sup_{x \in F(\theta_{k+1})} [u(\theta_{k+1}, x) - u(\theta_k, x)] \right\} \geq 0.$$

Since n is fixed and ϵ could be arbitrarily close to 0, we conclude that

$$\sum_{k=1}^{n} \left\{ \sup_{x \in F(\theta_{k+1})} [u(\theta_{k+1}, x) - u(\theta_k, x)] \right\} \geq 0,$$

which is exactly the cyclical monotonicity condition. The necessity part is shown.

Sufficiency. Define $N(\theta, \theta') := \sup_{x \in F(\theta')} [u(\theta', x) - u(\theta, x)]$ and $D(\theta, \theta') = -N(\theta, \theta')$. For each $\epsilon \in (0, 1]$ and each pair of types $\theta, \theta' \in \Theta$, select an outcome $x^{\epsilon}_{\theta,\theta'} \in F(\theta')$

such that

$$u(\theta', x^{\epsilon}_{\theta,\theta'}) - u(\theta, x^{\epsilon}_{\theta,\theta'}) + \epsilon \geq N(\theta, \theta').$$

We aim to construct a multi-mechanism $(g^{\epsilon}_{\theta}, t^{\epsilon}_{\theta})_{\theta\in\Theta,\epsilon\in(0,1]}$ that implements F. Let $g^{\epsilon}_{\theta}(\theta') = x^{\epsilon}_{\theta,\theta'}$. It remains to construct the transfer rules.

By cyclical monotonicity, for any nonempty finite sequence $\{\theta_1, ..., \theta_n\}$, and any sequence of numbers $\{\epsilon_1, ..., \epsilon_n\} \subseteq (0, 1]$, we have

$$\sum_{k=1}^{n} [u(\theta_{k+1}, x^{\epsilon_k}_{\theta_k,\theta_{k+1}}) - u(\theta_k, x^{\epsilon_k}_{\theta_k,\theta_{k+1}}) + \epsilon_k] \geq 0.$$

This implies that

$$\sum_{k=1}^{n-1} [u(\theta_k, x^{\epsilon_k}_{\theta_k,\theta_{k+1}}) - u(\theta_{k+1}, x^{\epsilon_k}_{\theta_k,\theta_{k+1}}) - \epsilon_k] \leq u(\theta_1, x^{\epsilon_n}_{\theta_n,\theta_1}) - u(\theta_n, x^{\epsilon_n}_{\theta_n,\theta_1}) + \epsilon_n. \quad \text{(B.5)}$$

Let $\mathcal{S}_n(\theta, \theta')$ be the collection of all sequences $\{\theta_1, ..., \theta_n\} \subseteq \Theta$ with $\theta_1 = \theta$ and $\theta_n = \theta'$. Let $\mathcal{G}_n$ be the collection of all sequences of numbers $\{\epsilon_1, ..., \epsilon_n\}$ with $\epsilon_k \in (0, 1]$ for each $k \in \{1, ..., n\}$. For any sequence of types $S = \{\theta_1, ..., \theta_n\}$ and any sequence of numbers $G = \{\epsilon_1, ..., \epsilon_{n-1}\} \in \mathcal{G}_{n-1}$, define

$$\alpha(S, G) = \sum_{k=1}^{n-1} [u(\theta_k, x^{\epsilon_k}_{\theta_k,\theta_{k+1}}) - u(\theta_{k+1}, x^{\epsilon_k}_{\theta_k,\theta_{k+1}}) - \epsilon_k].$$

Define

$$H^*(\theta, \theta') = \sup_{n\geq 2} \left\{ \sup_{S\in\mathcal{S}_n(\theta,\theta'), G\in\mathcal{G}_{n-1}} \alpha(S, G) \right\}.$$

By inequality (B.5), we know

$$H^*(\theta, \theta') \leq N(\theta', \theta) + 1 < +\infty.$$

Furthermore, for any $\theta, \theta', \theta'' \in \Theta$ and any $\epsilon \in (0, 1]$, consider the sequence of types $S_1 = \{\theta_1 = \theta', \theta_2, ..., \theta_n = \theta''\}$ with arbitrary $G_1 = \{\epsilon_1, ..., \epsilon_{n-1}\} \in \mathcal{G}_{n-1}$ and the sequence of types $S_2 = \{\theta, \theta_1 = \theta', \theta_2, ..., \theta_n = \theta''\}$ with arbitrary $G_2 = \{\epsilon, \epsilon_1, ..., \epsilon_{n-1}\} \in$

$\mathcal{G}_n$, then by inequality (B.5) and the definition of α and H^*, we have

$$u(\theta, x^\epsilon_{\theta,\theta'}) - u(\theta', x^\epsilon_{\theta,\theta'}) - \epsilon + \alpha(S_1, G_1) = \alpha(S_2, G_2) \leq H^*(\theta, \theta'')$$

Again by the definition of H^*, we take the supremum of the LHS of the above inequality over S_1, G_1 and n, which leads to

$$u(\theta, x^\epsilon_{\theta,\theta'}) - u(\theta', x^\epsilon_{\theta,\theta'}) + H^*(\theta', \theta'') \leq H^*(\theta, \theta'') + \epsilon. \tag{B.6}$$

Now we define the transfer schemes. Fix some $\theta^* \in \Theta$. Let $t^\epsilon_{\theta'}(\theta) = u(\theta, x^\epsilon_{\theta',\theta}) - H^*(\theta, \theta^*)$ for any $\theta, \theta' \in \Theta$ and $\epsilon \in (0, 1]$. Then

$$u(\theta, g^\epsilon_{\theta'}(\theta)) - t^\epsilon_{\theta'}(\theta) = u(\theta, x^\epsilon_{\theta',\theta}) - [u(\theta, x^\epsilon_{\theta',\theta}) - H^*(\theta, \theta^*)] = H^*(\theta, \theta^*).$$

For any $\theta, \theta' \in \Theta$ and $\epsilon \in (0, 1]$, under mechanism $(\Theta, g^\epsilon_\theta, t^\epsilon_\theta)$, the payoff of type θ agent by reporting θ' is

$$\begin{aligned} u(\theta, g^\epsilon_\theta(\theta')) - t^\epsilon_\theta(\theta') &= u(\theta, x^\epsilon_{\theta,\theta'}) - [u(\theta', x^\epsilon_{\theta,\theta'}) - H^*(\theta', \theta^*)] \\ &\leq H^*(\theta, \theta^*) + \epsilon. \end{aligned}$$

The last inequality is by (B.6). For any randomized misreporting, the multi-mechanism $(g^\epsilon_\theta, t^\epsilon_\theta)$ bounds the deviating gain by ϵ. Since the multi-mechanism $(g^\epsilon_\theta, t^\epsilon_\theta)_{\theta\in\Theta,\epsilon\in(0,1]}$ includes all the single mechanisms indexed by $\epsilon \in (0, 1]$, truth-telling is optimal by the maxmin criterion. Obviously, $H^*(\theta, \theta^*)$ is bounded below and thus non-triviality holds. Consistency is ensured by the fact that $x^\epsilon_{\theta,\theta'} \in F(\theta')$. The sufficiency is proved.

Weakly Undominated Condition. In the above proof, we fix $\theta^* \in \Theta$ and construct the multi-mechanism $(g^\epsilon_\theta, t^\epsilon_\theta)_{\theta\in\Theta,\epsilon\in(0,1]}$ such that for each $\theta, \theta' \in \Theta$ and $\epsilon \in (0, 1]$,

$$\begin{aligned} u(\theta, g^\epsilon_{\theta'}(\theta)) - t^\epsilon_{\theta'}(\theta) &= H^*(\theta, \theta^*), \\ u(\theta, g^\epsilon_\theta(\theta')) - t^\epsilon_\theta(\theta') &\leq H^*(\theta, \theta^*) + \epsilon. \end{aligned}$$

Now we construct a new multi-mechanism $(g^\epsilon_\theta, \hat{t}^\epsilon_\theta)_{\theta\in\Theta,\epsilon\in(0,1]}$ to implement the choice correspondence by slightly modifying the transfer schemes such that for each $\theta, \theta' \in \Theta$ and $\epsilon \in (0,1]$, $\hat{t}^\epsilon_\theta(\theta) = t^\epsilon_\theta(\theta) - 2\epsilon$ and $\hat{t}^\epsilon_\theta(\theta') = t^\epsilon_\theta(\theta')$ for all $\theta' \neq \theta$. To see that $(g^\epsilon_\theta, \hat{t}^\epsilon_\theta)_{\theta\in\Theta,\epsilon\in(0,1]}$ implements the choice correspondence, note that the truth-telling condition continues to hold as well as the other two conditions. For type θ agent, truth-telling is strictly better than any misreporting in the single-mechanism $(g^\epsilon_\theta, \hat{t}^\epsilon_\theta)$. Hence, Theorem 2.3.2 goes through with the additional weakly undominated condition.

B.3 Proof of Theorem 2.4.3

The necessity part is clear. For the sufficiency part, since F^a is implementable under $\bar{u}^a$ for each $a \in A$, we can consider a multi-mechanism $(\bar{g}^a_i, \bar{t}^a_i)_{i\in I^a}$ implementing F^a for each $a \in A$ such that $\bar{g}^a_i$ maps each reported type of agent a to a deterministic outcome. Now, we construct a multi-mechanism $(g_i, t_i)_{i\in\cup_{a\in A} I^a}$ to implement F based on $\{(\bar{g}^a_i, \bar{t}^a_i)_{i\in I^a}\}_{a\in A}$.

First, for any given $a \in A$, $\theta^a \in \Theta^a$ and $i \in I^a$, denote $\bar{g}^a_i(\theta^a)$ as γ^{-a}. We know that $\gamma^{-a} \in F^a(\theta^a)$ and thus we can define $g_i(\theta^a, \theta^{-a}) = \gamma^{-a}(\theta^{-a})$ and $t^a_i(\theta^a, \theta^{-a}) = \bar{t}^a_i(\theta^a)$ for any $\theta^{-a} \in \Theta^{-a}$. Then it remains to construct $t^{a'}_i(\theta)$ with $a' \neq a$ and $i \in I^a$. Let $t^{a'}_i(\theta) = K_i(\theta^{a'})$ for all $a' \neq a$ with $i \in I^a$ and $a, a' \in A$, such that for any $a \in A$, any $\theta^a \in \Theta^a$ and any $i \in \cup_{a'\neq a} I^{a'}$,

$$\inf_{j\in I^a}[\bar{u}^a(\theta^a, \bar{g}^a_j(\theta^a)) - \bar{t}^a_j(\theta^a)] \leq \int_{\Theta^{-a}} u^a(\theta^a, \theta^{-a}, g_i(\theta^a, \theta^{-a}))P^{-a}[d\theta^{-a}] - K_i(\theta^a).$$

Such $K_i(\theta^a)$ exists due to our boundedness assumption. One can immediately verify that under the multi-mechanism $(g_i, t_i)_{i\in\cup_{a\in A} I^a}$, when other agents are always telling the truth, each agent a, when reporting the true type, receives the same infimum payoff as in the multi-mechanism $(\bar{g}^a_i, \bar{t}^a_i)_{i\in I^a}$. Moreover, misreporting yields agent a a lower payoff under the multi-mechanism $(g_i, t_i)_{i\in\cup_{a\in A} I^a}$ than that under $(\bar{g}^a_i, \bar{t}^a_i)_{i\in I^a}$ if other agents are always telling the truth. This is due to the fact that $(g_i, t_i)_{i\in\cup_{a\in A} I^a}$

can be regarded an expansion of $(\bar{g}_i^a, \bar{t}_i^a)_{i\in I^a}$ and that agents are using the maxmin criterion. Thus, truth-telling constitutes an equilibrium, and $(g_i, t_i)_{i\in\cup_{a\in A}I^a}$ indeed implements the choice correspondence F.

B.4 Proof of Proposition 2.5.5

To show that F satisfies cyclical monotonicity, it is without loss of generality to consider any sequence $\{\theta^1, ..., \theta^n\}$[1] with $\theta^i \neq \theta^j$ for all $i \neq j$.[2] Since $\succsim^t$ is a total order, for each $i \neq j$, either $\theta^i \succ^t \theta^j$ or $\theta^j \succ^t \theta^i$. Consider the selection $\bar{x}_\theta$ and $\underline{x}_\theta$ for each type θ that satisfies the requirement of binary monotonicity. Let $N = \{1, ..., n\}$, and we want to show that

$$\sum_{k\in N:\theta^{k+1}\succ^t\theta^k} [u(\theta^{k+1}, \bar{x}_{\theta^{k+1}}) - u(\theta^k, \bar{x}_{\theta^{k+1}})] + \sum_{k\in N:\theta^k\succ^t\theta^{k+1}} [u(\theta^{k+1}, \underline{x}_{\theta^{k+1}}) - u(\theta^k, \underline{x}_{\theta^{k+1}})] \geq 0, \tag{B.7}$$

which proves the proposition. Fixing the sequence $\{\theta^1, ..., \theta^n\}$, define the sequence $\{\theta_1, ..., \theta_n\}$ such that $\{\theta^1, ..., \theta^n\} = \{\theta_1, ..., \theta_n\}$ and $\theta_k \succ^t \theta_{k-1}$ for $k \in \{2, ..., n\}$. Essentially, there is a permutation $\pi : \{1, ..., n\} \to \{1, ..., n\}$ such that $\theta^k = \theta_{\pi(k)}$ for each $k \in \{1, ..., n\}$. Consider any k with $\theta^k = \theta_m$ and $\theta^{k+1} = \theta_s$. If $m < s$, we have $\theta^{k+1} \succ^t \theta^k$ and

$$\begin{aligned} &u(\theta^{k+1}, \bar{x}_{\theta^{k+1}}) - u(\theta^k, \bar{x}_{\theta^{k+1}}) \\ &= \sum_{r=m}^{s-1} [u(\theta_{r+1}, \bar{x}_{\theta^{k+1}}) - u(\theta_r, \bar{x}_{\theta^{k+1}})]. \end{aligned}$$

Similarly, if $m > s$, we have $\theta^k \succ^t \theta^{k+1}$ and

$$\begin{aligned} &u(\theta^{k+1}, \underline{x}_{\theta^{k+1}}) - u(\theta^k, \underline{x}_{\theta^{k+1}}) \\ &= \sum_{r=s}^{m-1} [u(\theta_r, \underline{x}_{\theta^{k+1}}) - u(\theta_{r+1}, \underline{x}_{\theta^{k+1}})]. \end{aligned}$$

[1] Here we use superscripts to differentiate different types instead of subscripts. Subscripts will be used later.

[2] A sequence of types with repetitions can be divided into multiple subsequences without repetitions. When cyclical monotonicity holds for each subsequence, it also holds for the original sequence.

Without loss of generality, we can assume that $\pi(1) = 1$. Then by our assumption, $\theta^n \succ^t \theta^1$. We call θ^k is locally $\succsim^t$-maximal if $\theta^k \succ \theta^{k+1}$ and $\theta^k \succ \theta^{k-1}$ where $\theta^0 = \theta^n$ and $\theta^{n+1} = \theta^1$. A locally $\succsim^t$-minimal type can be defined similarly. Then there exists a subsequence of locally $\succsim^t$-maximal and $\succsim^t$-minimal types as $\{\theta^{l_1}, ..., \theta^{l_{2m}}\}$ with $1 = l_1 < l_2 < \cdots < l_{2m} \leq n$. θ^{l_p} is a locally $\succsim^t$-minimal type if p is odd and a locally $\succsim^t$-maximal type if p is even. This implies that for each $p = 1, ..., m$, $\pi(k)$ is strictly increasing for $l_{2p-1} \leq k \leq l_{2p}$ and strictly decreasing for $l_{2p} \leq k \leq l_{2p+1}$ where $l_{2m+1} := n + 1$. We have

$$
\begin{aligned}
&\sum_{k\in N:\theta^{k+1}\succ^t\theta^k} [u(\theta^{k+1}, \bar{x}_{\theta^{k+1}}) - u(\theta^k, \bar{x}_{\theta^{k+1}})] \\
&+ \sum_{k\in N:\theta^{k}\succ^t\theta^{k+1}} [u(\theta^{k+1}, \underline{x}_{\theta^{k+1}}) - u(\theta^k, \underline{x}_{\theta^{k+1}})] \\
&= \sum_{p=1}^{m} \sum_{k=l_{2p-1}}^{l_{2p}-1} [u(\theta^{k+1}, \bar{x}_{\theta^{k+1}}) - u(\theta^k, \bar{x}_{\theta^{k+1}})] \\
&+ \sum_{p=1}^{m} \sum_{k=l_{2p}}^{l_{2p+1}-1} [u(\theta^{k+1}, \underline{x}_{\theta^{k+1}}) - u(\theta^k, \underline{x}_{\theta^{k+1}})] \\
&= \sum_{p=1}^{m} \sum_{k=l_{2p-1}}^{l_{2p}-1} \sum_{a=\pi(k)}^{\pi(k+1)-1} [u(\theta_{a+1}, \bar{x}_{\theta^{k+1}}) - u(\theta_a, \bar{x}_{\theta^{k+1}})] \\
&+ \sum_{p=1}^{m} \sum_{k=l_{2p}}^{l_{2p+1}-1} \sum_{a=\pi(k+1)}^{\pi(k)-1} [u(\theta_{a+1}, \underline{x}_{\theta^{k+1}}) - u(\theta_a, \underline{x}_{\theta^{k+1}})].
\end{aligned}
$$

The summation is now divided into two parts. Note that the sequence starts and ends at θ_1. For any component in the first part as $u(\theta_{a+1}, \bar{x}_{\theta^{k+1}}) - u(\theta_a, \bar{x}_{\theta^{k+1}})$ for some $p \in \{1, ..., m\}$, $l_{2p-1} \leq k \leq l_{2p} - 1$ and $\pi(k) \leq a \leq \pi(k+1) - 1$, we can choose a component in the second part as $u(\theta_{a+1}, \underline{x}_{\theta^{k'+1}}) - u(\theta_a, \underline{x}_{\theta^{k'+1}})$ for some $p' \geq p$, $l_{2p'} \leq k' \leq l_{2p'+1} - 1$ with the same a. Such a mapping can be constructed to be well-defined and one-to-one. As a result, to show that inequality (B.7) holds, it suffices to prove for all possible p, k, a and the corresponding p', k',

$$
u(\theta_{a+1}, \bar{x}_{\theta^{k+1}}) - u(\theta_a, \bar{x}_{\theta^{k+1}}) \geq u(\theta_{a+1}, \underline{x}_{\theta^{k'+1}}) - u(\theta_a, \underline{x}_{\theta^{k'+1}}). \tag{B.8}
$$

Recall that $\theta^i \neq \theta^{i+1}$ for all $i = 1, ..., n$. Also, note that $\theta^{k+1} \succ^t \theta_a \succsim^t \theta^{k'+1}$. By binary monotonicity, we have $\bar{x}_{\theta^{k+1}} \trianglerighteq \underline{x}_{\theta^{k'+1}}$. Moreover, since $\theta_{a+1} \succ^t \theta_a$ and u exhibits increasing differences, inequality (B.8) holds. This completes the proof.

B.5 Proof of Proposition 2.5.6

We only need to show that weak monotonicity implies binary monotonicity. For any two types θ and θ' with $\theta \succ^t \theta'$, by weak monotonicity, we have

$$u(\theta, \bar{x}_\theta) - u(\theta', \bar{x}_\theta) + u(\theta', \underline{x}_{\theta'}) - u(\theta, \underline{x}_{\theta'}) \geq 0.$$

By strictly increasing differences and completeness of $\trianglerighteq$, we must have $\bar{x}_\theta \trianglerighteq \underline{x}_{\theta'}$. This finishes the proof.

B.6 Proof of Theorem 2.6.2

By Theorem 2.4.3, F is implementable in the multi-agent case if and only if F^a is implementable in the single-agent case for each $a \in A$, where the definition of F^a is given in Section 2.4. For any $\gamma^{-a} \in A^{\Theta^{-a}}$, agent a only cares about whether $\gamma^{-a}(\theta^{-a})$ equals a or not, which specifies whether he wins the project. Thus, for each γ^{-a}, define

$$W^a(\gamma^{-a}) := \sum_{\theta^{-a}:\gamma^{-a}(\theta^{-a})=a} P^{-a}[\theta^{-a}]$$

as the interim winning probability of agent a conditional on the chosen outcome γ^{-a}. Agent a's payoff (without transfer) under $\gamma^{-a} \in A^{\Theta^{-a}}$ is given by

$$\bar{u}^a(\theta^a, \gamma^{-a}) = v^a(\theta^a) W^a(\gamma^{-a}).$$

Note that W^a induces a weak order $\succsim^a$ over agent a's outcome space $A^{\Theta^{-a}}$ such that $\gamma^{-a} \succsim^a \hat{\gamma}^{-a}$ if and only if $W(\gamma^{-a}) \geq W(\hat{\gamma}^{-a})$. Moreover, since we assume that $v^a(\theta^a) \neq v^a(\hat{\theta}^a)$ for each $\theta^a \neq \hat{\theta}^a$, v^a induces a total order $\succsim^{a,t}$ over agent a's type

space such that $\theta^a \succsim^{a,t} \hat{\theta}^a$ if and only if $v^a(\theta^a) \geq v^a(\hat{\theta}^a)$. By Section 2.5, $\bar{u}^a$ satisfies strictly increasing differences with respect to $\succsim^a$ and $\succsim^{a,t}$. Therefore, by Proposition 2.5.6, any choice correspondence $F^a : \Theta^a \rightrightarrows A^{\Theta^{-a}}$ is implementable under $\bar{u}^a$ if and only if $v^a(\theta^a) > v^a(\hat{\theta}^a)$ implies

$$\sup_{\gamma^{-a} \in F^a(\theta^a)} W^a(\gamma^{-a}) \geq \inf_{\hat{\gamma}^{-a} \in F^a(\hat{\theta}^a)} W^a(\hat{\gamma}^{-a}).$$

Note that F^a is induced by F, where

$$F(\theta) = \left\{ a \in A : \theta^b \not\succ_G \theta^a, \forall b \in A \right\},$$

$$F^a(\theta^a) = \left\{ \gamma^{-a} \in A^{\Theta^{-a}} : \gamma^{-a}(\theta^{-a}) \in F(\theta^a, \theta^{-a}) \right\}.$$

Simple calculation indicates that

$$\sup_{\gamma^{-a} \in F^a(\theta^a)} W^a(\gamma^{-a}) = \mathcal{H}^a_{\succsim_G}(\theta^a),$$

$$\inf_{\hat{\gamma}^{-a} \in F^a(\hat{\theta}^a)} W^a(\hat{\gamma}^{-a}) = \mathcal{L}^a_{\succsim_G}(\hat{\theta}^a).$$

By Theorem 2.4.3, the implementability of F^a for each $a \in A$ is equivalent to the implementability of F. The proof is finished.

B.7 Proof of Proposition 2.7.4

The necessity part is clear. For sufficiency, let $(g_i, t_i)_{i \in I}$ implement F. Define transfer schemes $\{\hat{t}_i\}_{i \in I}$ such that for each $i \in I$, $\theta^a \in \Theta^a$ and each $\theta^{-a} \in \Theta^{-a}$,

$$\begin{aligned} &\int_X u^a(\theta^a, \theta^{-a}, x) g_i(\theta^a, \theta^{-a})[dx] - \hat{t}_i^a(\theta^a, \theta^{-a}) \\ =& \int_{\Theta^{-a}} \left(\int_X u^a(\theta^a, \hat{\theta}^{-a}, x) g_i(\theta^a, \hat{\theta}^{-a})[dx] - t_i^a(\theta^a, \hat{\theta}^{-a}) \right) P^{-a}[d\hat{\theta}^{-a}]. \end{aligned}$$

We argue that $(g_i, \hat{t}_i)_{i\in I}$ η-implements F. Note that for any $\theta^a \in \Theta^a$,

$$\begin{aligned}
&\inf_{i\in I}\left\{\int_{\Theta^{-a}}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\theta^a,\theta^{-a})[dx]-t_i^a(\theta^a,\theta^{-a})\right)P^{-a}[d\theta^{-a}]\right\}\\
=&\inf_{i\in I}\left\{\int_{\Theta^{-a}}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\theta^a,\theta^{-a})[dx]-\hat{t}_i^a(\theta^a,\theta^{-a})\right)P^{-a}[d\theta^{-a}]\right\}\\
=&\int_{\Theta^{-a}}\inf_{i\in I}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\theta^a,\theta^{-a})[dx]-\hat{t}_i^a(\theta^a,\theta^{-a})\right)P^{-a}[d\theta^{-a}]\\
>&-\infty,
\end{aligned}$$

and for any $\beta \in \Delta(\Theta^a)$,

$$\begin{aligned}
&\inf_{i\in I}\left\{\int_{\Theta^{-a}}\int_{\Theta^a}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\hat{\theta}^a,\theta^{-a})[dx]-t_i^a(\hat{\theta}^a,\theta^{-a})\right)\beta[d\hat{\theta}^a]P^{-a}[d\theta^{-a}]\right\}\\
=&\inf_{i\in I}\left\{\int_{\Theta^{-a}}\int_{\Theta^a}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\hat{\theta}^a,\theta^{-a})[dx]-\hat{t}_i^a(\hat{\theta}^a,\theta^{-a})\right)\beta[d\hat{\theta}^a]P^{-a}[d\theta^{-a}]\right\}\\
\geq&\int_{\Theta^{-a}}\inf_{i\in I}\left(\int_{\Theta^a}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\hat{\theta}^a,\theta^{-a})[dx]-\hat{t}_i^a(\hat{\theta}^a,\theta^{-a})\right)\beta[d\hat{\theta}^a]\right)P^{-a}[d\theta^{-a}].
\end{aligned}$$

Hence, for any $\theta^a \in \Theta^a$ and any $\beta \in \Delta(\Theta^a)$,

$$\begin{aligned}
&\int_{\Theta^{-a}}\inf_{i\in I}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\theta^a,\theta^{-a})[dx]-\hat{t}_i^a(\theta^a,\theta^{-a})\right)P^{-a}[d\theta^{-a}]\\
\geq&\int_{\Theta^{-a}}\inf_{i\in I}\left(\int_{\Theta^a}\left(\int_X u^a(\theta^a,\theta^{-a},x)g_i(\hat{\theta}^a,\theta^{-a})[dx]-\hat{t}_i^a(\hat{\theta}^a,\theta^{-a})\right)\beta[d\hat{\theta}^a]\right)P^{-a}[d\theta^{-a}].
\end{aligned}$$

We are done.

B.8 Proof of Proposition 2.7.5

Let $J = \cup_{\theta\in\Theta} I_\theta$ where each I_θ is a replicate of I such that each index $i \in I$ corresponds to the index $i_\theta \in I_\theta$. For each $i_\theta \in J$, let $\bar{g}_{i_\theta} = g_i$. Let $\hat{t}_{i_\theta}(\theta) = t_i(\theta)$ and $\hat{t}_{i_\theta}(\theta') = K_{i,\theta}$ for all $\theta' \neq \theta$. We let $K_{i,\theta} < t_i(\theta')$ for all $\theta' \neq \theta$. This is feasible since the transfers are assumed to be uniformly bounded. Obviously, the new multi-mechanism $(\bar{g}_j, \hat{t}_j)_{j\in J}$ still implements the choice correspondence. We can pick $K_{i,\theta}$ such that each single-mechanism yields the same expected transfer. Finally, we can

choose a constant C and let $\bar{t}_j(\theta) = \hat{t}_j(\theta) - C$ for each j and each θ such that

$$\inf_{i \in I} \left(\int_{\theta \in \Theta} t_i(\theta) P[d\theta] \right) = \int_{\theta \in \Theta} \bar{t}_j(\theta) P[d\theta], \ \forall j \in J.$$

B.9 Proof of Proposition 2.7.6

Since we can always add a constant transfer to each single-mechanism, we just need to show that the multi-mechanism we construct satisfies that $\sum_{a \in A} \int_{\theta \in \Theta} \bar{t}_j^a(\theta) P[d\theta]$ is a constant for each $j \in J$.

Let $J = \cup_{a \in A} I^a$ where each I^a is a replicate of I such that each $i \in I$ corresponds to $i^a \in I^a$. Let $\bar{g}_{i^a} = g_i$ for each $i^a \in J$. Let $\bar{t}_{i^a}^a(\theta) = t_i^a(\theta)$ for each $i^a \in J$ and $\theta \in \Theta$ and $\bar{t}_{i^a}^{a'}(\theta) = K_{i,a}$ for each $i^a \in J$, each $a' \neq a$ and each $\theta \in \Theta$. By letting $K_{i,a} < t_i^{a'}(\theta)$ for each $i \in I$ and $a \in A$ and $a' \in A$, we can ensure that the new multi-mechanism $(\bar{g}_j, \bar{t}_j)_{j \in J}$ implements the choice correspondence. Moreover, we can pick $K_{i,a}$ to ensure that $\sum_{a \in A} \int_{\theta \in \Theta} \bar{t}_j^a(\theta) P[d\theta]$ is a constant for each $j \in J$. We are done.

Appendix C

Appendices for Chapter 3

C.1 Proof of Lemma 3.3.4

By the form of the objective function, it is easy to see that $\hat{\rho}^*$ solves either $\max_{\hat{\rho}\in[\underline{\rho},\overline{\rho}]^\infty} \mathbb{E}[\theta|a,\hat{\rho}]$ or $\min_{\hat{\rho}\in[\underline{\rho},\overline{\rho}]^\infty} \mathbb{E}[\theta|a,\hat{\rho}]$.

Given a distribution of observables, H with density h, recall that

$$v(\hat{\rho}) \equiv \mathbb{E}[\theta|a,\hat{\rho}] = \int \frac{\hat{\rho}(x)\hat{s}\left(x,\hat{\rho}(x)\right)}{\int \hat{\rho}(x)h(x)dx} h(x)dx.$$

Fix a value $M \in [\underline{\rho},\overline{\rho}]$ and consider the problem:

$$\max_{\hat{\rho}[\underline{\rho},\overline{\rho}]^\infty} \{v(\hat{\rho}) : \int \hat{\rho}(x)h(x)dx = M\}$$
$$= \frac{1}{M} \max_{\hat{\rho}[\underline{\rho},\overline{\rho}]^\infty} \{\int \hat{\rho}(x)\hat{s}\left(x,\hat{\rho}(x)\right) h(x)dx : \int \hat{\rho}(x)h(x)dx = M\}.$$

where the last equality is justified because we are equating the denominator of v to M. By Lagrange multiplier Theorem in Banach spaces, we obtain that there is λ such that, for each x:

$$\hat{\rho}(x) \in \arg\max_{\rho\in[\underline{\rho},\overline{\rho}]} \{\rho\hat{s}(x,\rho) - \lambda(\rho - M)\}.$$

By the supermodularity in Assumption 3.3.3, we know the objective function of each of these optimizations is supermodular, so $\hat{\rho}(x)$ is increasing with x, according to Topkis' lemma. By affinity, the solution can be assumed to be an extreme point of the interval $[\underline{\rho}, \overline{\rho}]$. Therefore, for each M the solution is a threshold. Thus, maximizing over M's the solution must also be a threshold. Clearly, the same result holds for minimization and the proof is concluded.

C.2 Proof of Theorem 3.3.5

For any realization of observables a^N, let $F^N \in \Delta(\mathbb{R})$ be the empirical distribution of observables. We abuse notation to write $s^a(a^N, \hat{\rho}^N)$ as the vector in which the i-th entry is $s^a(a_i^N, \hat{\rho}_i^N)$. Given a conjecture $\hat{\rho}^N$, we know the backtracked signals $s^a(a_i^N, \hat{\rho}_i^N)$ are jointly normal with the state, allowing us to calculate the posterior mean as:

$$\mathbb{E}[\theta | a^N, \hat{\rho}^N] = \frac{\hat{\rho}^N \cdot s^a(a^N, \hat{\rho}^N) + \rho_\mu \mu}{\hat{\rho}^N \cdot \mathbb{1} + \rho_\mu}.$$

Define:

$$\underline{m}^N \equiv \min_{\hat{\rho} \in [\underline{\rho}, \overline{\rho}]^N} \mathbb{E}[\theta | a^N, \hat{\rho}^N],$$

and

$$\overline{m}^N \equiv \max_{\hat{\rho} \in [\underline{\rho}, \overline{\rho}]^N} \mathbb{E}[\theta | a^N, \hat{\rho}^N].$$

$\underline{m}^N$ and $\overline{m}^N$ are (random) bounds on posterior means. Assume that $\underline{\rho}^N$ and $\overline{\rho}^N$ are the respective maximizers.

Let $\hat{\rho} : \mathbb{R} \to [\underline{\rho}, \overline{\rho}]$ be a precision assignment. Let H be the real distribution of observables. Again, given a precision assignment, signals are jointly normally distributed with the state, so we can write the posterior mean as:

$$\mathbb{E}[\theta | \hat{\rho}] = \int \frac{\hat{\rho}(x) s^a(x, \hat{\rho}(x)) dF(x)}{\int \hat{\rho}(x) dF(x)}.$$

Finally, let:

$$\underline{m} = \min_{\hat{\rho}:\mathbb{R}\to[\underline{\rho},\overline{\rho}]} \mathbb{E}[\theta|\hat{\rho}] \qquad \overline{m} = \min_{\hat{\rho}:\mathbb{R}\to[\underline{\rho},\overline{\rho}]} \mathbb{E}[\theta|\hat{\rho}].$$

We start the proof by showing, in **Step 1**, that the random bounds on posterior means converge to $\underline{m}$ and $\overline{m}$ asymptotically. Then, we show that the latter are indeed asymptotic bounds of posterior means, proving part 1 of the theorem in **Step 2**. Finally, we show part 2 of the theorem in **Step 3**.

Step 1. $\underline{m}^N \xrightarrow{a.s.} \underline{m}$ **and** $\overline{m}^N \xrightarrow{a.s.} \overline{m}$**.**

Step 1.1. $\overline{m} = \frac{\underline{\rho}\int_{-\infty}^{\overline{m}} \boldsymbol{s}^{\boldsymbol{a}}(x,\underline{\rho})dF(x)+\overline{\rho}\int_{\overline{m}}^{\infty} \boldsymbol{s}^{\boldsymbol{a}}(x,\overline{\rho})dF(x)}{\underline{\rho}F(\overline{m})+\overline{\rho}(1-F(\overline{m}))}$.

By the proof of Lemma 3.3.4, $\underline{m}$ is solved by a threshold strategy. We can then write the optimization that determines it as:

$$\overline{m} = \arg\max_{a\in\mathbb{R}} \{v(a)\},$$

where $v(a) = \frac{\underline{\rho}\int_{-\infty}^{a} \boldsymbol{s}^{\boldsymbol{a}}(x,\underline{\rho})dF(x)+\overline{\rho}\int_{a}^{\infty} \boldsymbol{s}^{\boldsymbol{a}}(x,\overline{\rho})dF(x)}{\underline{\rho}F(a)+\overline{\rho}(1-F(a))}$.

The first order condition leads to:

$$\overline{m} = \frac{\underline{\rho}\int_{-\infty}^{\overline{m}} \boldsymbol{s}^{\boldsymbol{a}}(x,\underline{\rho})dF(x) + \overline{\rho}\int_{\overline{m}}^{\infty} \boldsymbol{s}^{\boldsymbol{a}}(x,\overline{\rho})dF(x)}{\underline{\rho}F(\overline{m}) + \overline{\rho}\,(1-F(\overline{m}))}$$

which implicitly defines the value $\overline{m}$ that solves that maximization problem. We show that the objective function is single-peaked, so that the first order condition is necessary and sufficient. The first derivative of v can be written as:

$$v'(a) = (v(a)-a)\,\frac{(\overline{\rho}-\underline{\rho})f(a)}{\underline{\rho}F(a) + \overline{\rho}\,(1-F(a))}.$$

First, notice that because the second term is positive for all $a \in \mathbb{R}$, the sign of v' is determined by $v(a) - a$. We claim that if there is $\underline{a}$ such that $v'(\underline{a}) > 0$, then $v'(a) > 0$ for all $a \leq \underline{a}$; similarly, if there is $\overline{a}$ such that $v'(\overline{a}) < 0$, then $v'(a) < 0$ for all $a \geq \overline{a}$. We prove the second, and the first follows by symmetry. Assume there is $\overline{a}$

such that $v'(\overline{a}) < 0$ and, to obtain a contradiction, let there be $a > \overline{a}$ with $v'(a) > 0$. Since v' is continuous, there must be $\overline{a} < b < a$ with $v'(b) = 0$, which implies $v(b) = b$. Choose the smallest such $b > \overline{a}$, so for $\overline{a} \leq x < b$, $v'(x) < 0$. We then have:

$$0 = v(b) - b < v(b) - \overline{a} = v(\overline{a}) - \overline{a} + \int_{\overline{a}}^{b} v'(x)dx < 0,$$

since $v'(\overline{a}) < 0$ implies $v(\overline{a}) < \overline{a}$. We have thus obtained a contradiction.

As $a \to -\infty$, $v(a) \to \int_{-\infty}^{\infty} xf(x)dx$, as all signals are assigned precision $\overline{\rho}$, leading to uniform weighting. Because we know F has a finite mean, that implies that we can find a sufficiently small number $\underline{a}$ such that $v(\underline{a}) - \underline{a} > 0$, implying $v'(\underline{a}) > 0$. Notice that the same should be true for all $a \leq \underline{a}$, so that v is an increasing function in $(-\infty, \underline{a}]$.

On the other hand, as $a \to \infty$, again we have $v(a) \to \int_{-\infty}^{\infty} xf(x)dx$, this time because all signals are receiving precision $\underline{\rho}$. Then, there is a sufficiently high number $\overline{a}$ with $v(a) - a < 0$, so $v'(a) < 0$ for all $a \geq \overline{a}$.

Hence v is single-peaked and the first order condition is necessary and sufficient. Because v' is continuous, there is $a^* \in [\underline{a}, \overline{a}]$ with $v'(a^*) = 0$, so the solution exists. We now prove uniqueness. Let a' satisfy $v'(a') = 0$, and let $a' > a^*$ without loss of generality. By the argument above, $v'(x) = 0$ for all $x \in [a^*, a']$. Then:

$$0 = v(a') - a' < v(a') - a^* = v(a^*) - a^* + \int_{a^*}^{a'} v'(x)dx = 0,$$

again, yielding a contradiction. Therefore a^* is unique. This concludes Step 1.1 By symmetry, we have the definition of $\underline{m}$.

Step 1.2. Approximating $\overline{m}^N$ using a threshold. In this step we show how to approximate the expectation $\overline{m}^N$ by the expectation generated by a threshold strategy as N grows large. For any realization of actions, a^N, let F^N be the associated empirical distribution of actions. We then define:

$$\tilde{m}^N = \max_{a \in \mathbb{R}} \frac{\underline{\rho} \int_{-\infty}^{a} \boldsymbol{s}^{\boldsymbol{a}}(x, \underline{\rho}) dF^N(x) + \overline{\rho} \int_{a}^{\infty} \boldsymbol{s}^{\boldsymbol{a}}(x, \overline{\rho}) dF^N(x)}{\underline{\rho} F^N(a) + \overline{\rho}\,(1 - F^N(a))}.$$

Call the objective function of the problem above $\Psi^N(a)$. At the same time, using the proof of Lemma 3.3.4 without assuming the distribution of observables is non-atomic, we obtain that $\overline{m}^N$ can be obtained by an assignment that is a threshold except for possibly one of the observables receiving an intermediate precision. Thus, we can find $\overline{m}^N$ through the alternative optimization:

$$\overline{m}^N = \max_{a,\rho\in[\underline{\rho},\overline{\rho}]} \left\{ \frac{\underline{\rho}\int_{-\infty}^{a-} \boldsymbol{s^a}(x,\underline{\rho})dF^N(x) + \rho \boldsymbol{s^a}(a,\rho)(F^N(a)-F^N(a-)) + \overline{\rho}\int_a^{\infty} \boldsymbol{s^a}(x,\overline{\rho})dF^N(x) + \frac{\rho_\mu}{N}\mu}{\underline{\rho}F^N(a-) + \rho(F^N(a)-F^N(a-)) + \overline{\rho}\,(1-F^N(a)) + \frac{\rho_\mu}{N}} \right\}.$$

We call the objective function above $\tilde{\Psi}^N(a,\rho)$. We next prove $\sup_{a\in\mathbb{R},\rho\in[\underline{\rho},\overline{\rho}]} |\tilde{\Psi}^N(a,\rho) - \Psi^N(a)| \xrightarrow{a.s} 0$. To see that, notice that for almost all sequences a, it must be that $\sup_a \left\{F^N(a) - F^N(a-)\right\} \leq \frac{1}{N}$. Applying that, the uniform convergence result immediately follows.

Denote

$$\Psi(a) = \frac{\underline{\rho}\int_{-\infty}^{a} \boldsymbol{s^a}(x,\underline{\rho})dF(x) + \overline{\rho}\int_a^{\infty} \boldsymbol{s^a}(x,\overline{\rho})dF(x)}{\underline{\rho}F(a) + \overline{\rho}\,(1-F(a))}.$$

where F is, again, the true distribution of observables.

Step 1.3. $\sup_{\boldsymbol{a}\in\mathbb{R}} |\Psi^N(\boldsymbol{a}) - \Psi(\boldsymbol{a})| \xrightarrow{a.s.} 0$ Given the Glivenko–Cantelli theorem, we know the empirical distribution function converges to the true cumulative distribution function uniformly over x, that is,

$$\|F^N - F\| := \sup_{x\in\mathbb{R}} |F^N(x) - F(x)| \xrightarrow{a.s.} 0.$$

For each real-valued function v, denote

$$F^N(v) = \int v dF^N \quad , \quad F(v) = \int v dF.$$

A class of real-valued functions $\mathcal{V}$ is defined to be a *P-Glivenko-Cantelli class of functions* if

$$\|F^N - F\|_{\mathcal{V}} := \sup_{v\in\mathcal{V}} |F^N(v) - F(v)| \xrightarrow{a.s.} 0.$$

Recall that the $L_1(F)$ norm is defined for real-valued functions such that

$$\|v\|_{L_1(F)} = \int |v| dF.$$

Given two real-valued functions l and u and $\epsilon > 0$, a *ε-bracket* $[l, u]$ is the set of all functions f such that $l \leq f \leq u$ and $\|u - l\|_{L_1(F)} \leq \varepsilon$. The *bracketing number* $N(\varepsilon, \mathcal{V}, \|\cdot\|_{L_1(F)})$ is the minimum number of ε-brackets needed to cover $\mathcal{V}$. The following theorem provides a sufficient condition for a P-Glivenko-Cantelli class.

Theorem C.2.1 *(Blum, 1955; DeHardt, 1971) If $N(\varepsilon, \mathcal{V}, \|\cdot\|_{L_1(F)}) < \infty$ for any $\varepsilon > 0$, then $\mathcal{V}$ is a P-Glivenko-Cantelli class.*

Denote

$$\mathcal{V}_1 = \Big\{ v_1^a : v_1^a(x) = \underline{\rho} \mathbb{1}_{\{x \leq a\}} + \overline{\rho} \mathbb{1}_{\{x > a\}}, \forall x \in \mathbb{R}, \text{ for some } a \in \mathbb{R} \Big\}.$$

$$\mathcal{V}_2 = \Big\{ v_2^a : v_2^a(x) = \underline{\rho} x \mathbb{1}_{\{x \leq a\}} + \overline{\rho} x \mathbb{1}_{\{x > a\}}, \forall x \in \mathbb{R}, \text{ for some } a \in \mathbb{R} \Big\}.$$

Easy to see

$$\Psi^N(a) = \frac{F^N(v_2^a)}{F^N(v_1^a)} \quad , \quad \Psi(a) = \frac{F(v_2^a)}{F(v_1^a)}.$$

Then we want to show that $\mathcal{V}_1$ and $\mathcal{V}_2$ are both P-Glivenko-Cantelli classes. Note that F is a continuous distribution whose expectation is well-defined, that is, $\int |x| dF < \infty$.

Fix $\varepsilon > 0$. For any $a > b$, the $L_1(F)$-distance between v_1^a and v_1^b is

$$\|v_1^a - v_1^b\|_{L_1(F)} = (\overline{\rho} - \underline{\rho}) \int_b^a dF(x).$$

Since $\int_{-\infty}^{\infty} dF(x) = 1$, for M large enough, we can find a finite increasing sequence $\{a_1, ..., a_M\}$ on the extended real line such that $a_1 = -\infty$, $a_M = \infty$ and

$$\int_{a_i}^{a_{i+1}} dF(x) = \frac{1}{M-1} \leq \frac{\varepsilon}{\overline{\rho} - \underline{\rho}}, \forall i = 1, ..., M-1$$

This is feasible as F is a continuous distribution. Then it is easy to show that the set of ε-brackets $\{[v_1^{a_i}, v_1^{a_{i+1}}] : i = 1, ..., M-1\}$ covers $\mathcal{V}_1$ and $N(\varepsilon, \mathcal{V}_1, \|\cdot\|_{L_1(F)}) \leq M - 1 < \infty$. Hence $\mathcal{V}_1$ is a P-Glivenko-Cantelli class.

Similarly, for any $a > b$, the $L_1(F)$-distance between v_2^a and v_2^b is

$$\|v_2^a - v_2^b\|_{L_1(P)} = (\overline{\rho} - \underline{\rho}) \int_b^a |x| dF(x).$$

Since $\int |x| dF < \infty$ and F is continuous, for M' large enough, again we can fine a finite increasing sequence $\{b_1, ..., b_{M'}\}$ on extended real line such that $b_1 = -\infty$, $b_{M'} = \infty$ and

$$\int_{b_i}^{b_{i+1}} |x| dF(x) = \frac{\int |x| dF}{M' - 1} \leq \frac{\varepsilon}{\overline{\rho} - \underline{\rho}}, \forall i = 1, ..., M' - 1.$$

Then it is easy to show that the set of ε-brackets $\{[v_2^{b_i}, v_2^{b_{i+1}}] : i = 1, ..., M'-1\}$ covers $\mathcal{F}_2$ and $N(\varepsilon, \mathcal{V}_2, \|\cdot\|_{L_1(F)}) \leq M' - 1 < \infty$. Hence $\mathcal{V}_2$ is a P-Glivenko-Cantelli class.

The definition of the P-Glivenko-Cantelli class implies that

$$\|F^N - F\|_{\mathcal{V}_1} = \sup_{v \in \mathcal{V}_1} |F^N(v) - F(v)| = \sup_{a \in \mathbb{R}} |F^N(v_1^a) - F(v_1^a)| \xrightarrow{a.s.} 0. \tag{C.1}$$

$$\|F^N - F\|_{\mathcal{V}_1} = \sup_{v \in \mathcal{V}_1} |F^N(v) - F(v)| = \sup_{a \in \mathbb{R}} |F^N(v_1^a) - F(v_1^a)| \xrightarrow{a.s.} 0. \tag{C.2}$$

Now we can show the convergence of Ψ^N.

$$\begin{aligned}
\sup_{a \in \mathbb{R}} |\Psi^N(a) - \Psi(a)| =& \sup_{a \in \mathbb{R}} |\frac{F^N(v_2^a)}{F^n(v_1^a)} - \frac{F(v_2^a)}{F(v_1^a)}| \\
\leq& \sup_{a \in \mathbb{R}} |\frac{F^N(v_2^a)}{F^N(v_1^a)} - \frac{F^N(v_2^a)}{F(v_1^a)}| + \sup_{a \in \mathbb{R}} |\frac{F^N(v_2^a)}{F(v_1^a)} - \frac{F(v_2^a)}{F(v_1^a)}| \\
\leq& \sup_{a \in \mathbb{R}} |\frac{F^N(v_2^a)}{F(v_1^a)F^N(v_1^a)}||F^N(v_1^a) - F(v_1^a)| \\
&+ \sup_{a \in \mathbb{R}} \frac{1}{|F(v_1^a)|} |F^N(v_2^a) - F(v_2^a)| \\
\leq& \sup_{a \in \mathbb{R}} |\frac{F^N(v_2^a)}{F(v_1^a)F^N(v_1^a)}| \sup_{a \in \mathbb{R}} |F^N(v_1^a) - F(v_1^a)| \\
&+ \sup_{a \in \mathbb{R}} \frac{1}{|F(v_1^a)|} \sup_{a \in \mathbb{R}} |F^N(v_2^a) - F(v_2^a)|.
\end{aligned}$$

Notice that $0 < \underline{\rho} \leq F(v_1^a) \leq \overline{\rho} < \infty$ and $0 < \underline{\rho} \leq F^N(v_1^a) \leq \overline{\rho} < \infty$ for each N. That is, $F(v_1^a)$ and $F^N(v_1^a)$ are uniformly bounded away from 0 and ∞. Also, by applying strong law of large numbers,

$$\sup_{a\in\mathbb{R}} |F^N(v_2^a)| \leq (\underline{\rho}+\overline{\rho}) \int |x| dF^N \xrightarrow{a.s.} (\underline{\rho}+\overline{\rho}) \int |x| dF < +\infty.$$

By equations C.1 and C.2, we know

$$\sup_{a\in\mathbb{R}} |\Psi^N(a) - \Psi(a)| \xrightarrow{a.s.} 0.$$

Step 1.4. $\overline{m}^N \xrightarrow{a.s.} \overline{m}$ This result follows directly from the following standard results about consistency of M- estimators. We include the proof for completeness.

Lemma C.2.2 *Suppose that*

1. $\sup_{a\in\mathbb{R},\rho\in[\underline{\rho},\overline{\rho}]} |\tilde{\Psi}^N(a,\rho) - \Psi(a)| \xrightarrow{a.s.} 0$,
2. $\overline{m}_N \in \arg\max_{a\in\mathbb{R},\rho\in[\underline{\rho},\overline{\rho}]} \tilde{\Psi}^N(a,\rho)$ *for each* N,
3. $\overline{m} = \arg\max_{a\in\mathbb{R}} \Psi(a)$ *is the unique maximum of* Ψ,

Then $\overline{m}_N \xrightarrow{a.s.} \overline{m}$.

Proof of Lemma C.2.2. We ignore the argument ρ throughout the proof without loss of generality. By conditions (2) and (3), we know $\tilde{\Psi}^N(\overline{m}_N) \geq \tilde{\Psi}^N(\overline{m})$ and $\Psi(\overline{m}) \geq \Psi(\overline{m}_N)$ for each N. Using these inequalities we have

$$\tilde{\Psi}^N(\overline{m}_N) - \Psi(\overline{m}_N) \geq \tilde{\Psi}^N(\overline{m}_N) - \Psi(\overline{m}) \geq \tilde{\Psi}^N(\overline{m}) - \Psi(\overline{m}).$$

Therefore from the above we have

$$|\tilde{\Psi}^N(\overline{m}_N) - \Psi(\overline{m})| \geq \max\left\{|\tilde{\Psi}^N(\overline{m}_N) - \Psi(\overline{m}_N)|, |\tilde{\Psi}^N(\overline{m}) - \Psi(\overline{m})|\right\} \geq \sup_{a\in\mathbb{R}} |\tilde{\Psi}^N(a) - \Psi(a)|.$$

Hence by condition (1), we know $|\tilde{\Psi}^N(\overline{m}_N) - \Psi(\overline{m})| \xrightarrow{a.s.} 0$. Finally, suppose by contradiction that $\overline{m}_N$ does not converge to $\overline{m}$ almost surely. Then there exists an

event M with positive probability such that for all $\omega \in M$, $\overline{m}_N(\omega) \not\to \overline{m}(\omega)$. As $\overline{m}$ is the unique minimum of Ψ by condition (3), $\tilde{\Psi}(\overline{m}_N(\omega)) \not\to \Psi(\overline{m}(\omega))$. Again condition (1) implies that $|\tilde{\Psi}^N(\overline{m}_N) - \Psi(\overline{m}_N)| \xrightarrow{a.s.} 0$. Hence we know that there exists $M' \subseteq M$ with positive probability such that for all $\omega \in M'$, $\tilde{\Psi}^N(\overline{m}_N(\omega)) \not\to \Psi(\overline{m}(\omega))$, which contradicts with $|\tilde{\Psi}^N(\overline{m}_N) - \Psi(\overline{m})| \xrightarrow{a.s.} 0$. Thus, we have $\overline{m}_N \xrightarrow{a.s.} \overline{m}$. ■

Now it suffices to show that the conditions in Lemma C.2.2 holds in our case. Condition (1) is shown in Step 1.2 and 1.3. Explicitly: $\sup_{a,\rho} |\tilde{\Psi}^N(a,\rho) - \Psi(a)| \xrightarrow{a.s} 0$ and $\sup_a |\Psi^N(a) - \Psi(a)| \xrightarrow{a.s.} 0$ imply that condition. Condition (2) holds by the definition of $\overline{m}_N$. Condition (3) is shown in the proof of Step 1.1. This completes the proof for $\overline{m}_N \xrightarrow{a.s.} \overline{m}$. The same arguments apply for showing $\underline{m}_N \xrightarrow{a.s.} \underline{m}$.

Step 2. Part 1 of Theorem — Boundedness of belief means. For any N, with observables a^N and conjectured precisions $\hat{\rho}^N$, recall we have:

$$\theta | s^N, \hat{\rho}^N \sim \left(\frac{\sum_{i=1}^N \hat{\rho}_i \boldsymbol{s}^{\boldsymbol{a}}(a_i, \hat{\rho}_i) + \rho_\mu \mu}{\sum_{i=1}^N \hat{\rho}_i + \rho_\mu}, \Big(1 - \frac{\sum_{i=1}^N \hat{\rho}_i}{\sum_{i=1}^N \hat{\rho}_i + \rho_\mu}\Big) \frac{1}{\rho_\mu} \right). \tag{C.3}$$

Since $\hat{\rho}_i \geq \underline{\rho} > 0$, it is clear that $\lim_{N\to\infty} \frac{\sum_{i=1}^N \hat{\rho}_i}{\sum_{i=1}^N \hat{\rho}_i + \rho_\mu} = 1$, so the variance converges to zero for all sequences of signal realizations.

As for the posterior mean, notice that, by definition of $\underline{m}^N$, $\overline{m}^N$:

$$\underline{m}^N \leq \frac{\sum_{i=1}^N \hat{\rho}_i \boldsymbol{s}^{\boldsymbol{a}}(a_i, \hat{\rho}_i) + \rho_\mu \mu}{\sum_{i=1}^N \hat{\rho}_i + \rho_\mu} \leq \overline{m}^N.$$

By taking limit inferior in the first inequality above and limit superior in the second, we obtain, using the result in Step 1, that for almost all sequences of signal realizations, the asymptotic bounds on expected values hold.

Step 3. Part 2 of Theorem — Limit Set of Posteriors Fix a sequence of realizations a. We want to characterize the set of distributions the posterior beliefs of the decision maker converge to, $\mathbb{P}_\infty(a)$. By C.3, it is clear that a necessary condition for weak convergence is that the posterior mean $\frac{\sum_{i=1}^N \hat{\rho}_i \boldsymbol{s}^{\boldsymbol{a}}(a_i, \hat{\rho}_i) + \rho_\mu \mu}{\sum_{i=1}^N \hat{\rho}_i + \rho_\mu}$

converges. We can then focus on sequences with converging means. Define $b = \lim_{N\to\infty} \frac{\sum_{i=1}^{N} \hat{\rho}_i \boldsymbol{s}^{\boldsymbol{a}}(a_i,\hat{\rho}_i)+\rho_\mu \mu}{\sum_{i=1}^{N} \hat{\rho}_i+\rho_\mu}$.

We can write the characteristic function of $P_N(s^N, \hat{\rho}^N)$ as:

$$\varphi^N(t) = e^{it\left\{ \frac{\sum_{i=1}^{N} \hat{\rho}_i \boldsymbol{s}^{\boldsymbol{a}}(a_i,\hat{\rho}_i)+\rho_\mu \mu}{\sum_{i=1}^{N} \hat{\rho}_i+\rho_\mu} - \frac{1}{2}\left(1-\frac{\sum_{i=1}^{N} \hat{\rho}_i}{\sum_{i=1}^{N} \hat{\rho}_i+\rho_\mu}\right)\frac{1}{\rho_\mu}\right\}}.$$

By Step 2, the variance converges to zero. We then have, for all t:

$$\varphi^N(t) \to e^{it\ b}.$$

which is the characteristic function of δ_b. Then, by Levy's continuity theorem: $P_N(s^N, \hat{\rho}^N) \xrightarrow{w} \delta_b$.

We finally show that any $b \in [\underline{m}, \overline{m}]$ can be achieved. For that, fix a threshold assignment $\rho : \mathbb{R} \to \{\underline{\rho}, \overline{\rho}\}$. Then $\{\rho(a_i)\boldsymbol{s}^{\boldsymbol{a}}(a_i, \rho(a_i))\}_{i=1,\dots}$ is a sequence of independent signals with uniformly bounded variance. Then, by the strong law of large numbers:

$$\begin{aligned}\frac{\sum_{i=1}^{N} \rho(s_i)\boldsymbol{s}^{\boldsymbol{a}}(a_i,\rho(a_i)) + \rho_\mu \mu}{\sum_{i=1}^{N} \rho(s_i) + \rho_\mu} &= \frac{N \int \rho(x)\boldsymbol{s}^{\boldsymbol{a}}(x,\rho(x))dF^N(x) + \rho_\mu \mu}{N \int \rho(x)dF^N(x) + \rho_\mu} \\ &\xrightarrow{a.s.} \frac{\int \rho(x)\boldsymbol{s}^{\boldsymbol{a}}(x,\rho(x))dF(x)}{\int \rho(x)dF(x)}.\end{aligned}$$

We finish this step by showing that by appropriately choosing the function ρ, $\frac{\int \rho(x)sdF(x)}{\int \rho(x)dF(x)}$ can achieve any point between $\underline{m}$ and $\overline{m}$. To see that, recall that $\overline{m} = \max_a \Psi(a)$. It should be clear that $\mu = \min_a \Psi(a)$. Since Ψ is continuous, by choosing different a's any number in $[\mu, \overline{m}]$ can be achieved. Because any a corresponds to a particular threshold assignment ρ, this means that $\frac{\int \rho(x)\boldsymbol{s}^{\boldsymbol{a}}(x,\rho(x))dF(x)}{\int \rho(x)dF(x)}$ can achieve any value in $[\mu, \overline{m}]$. With the symmetric argument for $\underline{m}$, we obtain the result and complete the proof.

C.3 Proof of Theorem 3.4.1

Define

$$\Gamma^N(g) \equiv \max_{p \in \mathbb{P}^{\boldsymbol{\alpha}}(a^N)} \mathbb{E}_p \left[u(g - \theta) \right].$$

By definition, assuming that the limits exist, we have:

$$\lim_{N \to \infty} g^*(s^N) = \lim_{N \to \infty} \arg\min_g \Gamma^N(g).$$

Also denote

$$\Gamma(g) = \max\{u(g - \overline{m}), u(g - \underline{m})\}$$

where $\underline{m}$ and $\overline{m}$ are defined in Theorem 3.3.5.

We start with introducing an auxiliary problem with finitely many signals by ignoring the effect of any moment of the posterior distribution that is not the mean. Explicitly:

$$\tilde{\Gamma}^N(g) \equiv \max_{p \in \mathbb{P}^{\boldsymbol{\alpha}}(a^N)} u(g - \mathbb{E}_p[\theta]) = \max \left\{ u(g - \overline{m}), u(g - \underline{m}) \right\},$$

where $\overline{m}$ and $\underline{m}$ are defined in Theorem 3.3.5 and the equality follows from the fact that u is convex.

The result of the proposition is a consequence of the following lemma.

Lemma C.3.1 *Let f^N be a sequence of random mappings such that $x^N \in \arg\min_{x \in \mathbb{R}} f^N(x)$, for all $N \in \mathbb{N}$. Assume there is another random mapping f and that the following are satisfied:*

1. $\sup_{x \in C} |f(x) - f^N(x)| \xrightarrow{a.s} 0$, *as $N \to \infty$, for all compact sets $C \subset \mathbb{R}$.*

2. $x^* \in \arg\min_{x \in \mathbb{R}} f(x)$ *is the unique minimum of f.*

3. *The sequence x^N is uniformly bounded almost everywhere.*

Then $x^N \xrightarrow{a.s} x^*$.

Proof of Lemma C.3.1. By condition (3), there exists an event M with $\mathbb{P}(M) = 1$ such that for all $\omega \in M$, there is a compact set $C(\omega) \subseteq \mathbb{R}$ with $\{x^N(\omega)\}_{N\geq 1} \cup \{x^*(\omega)\} \subseteq C(\omega)$. By condition (1), we can find $M' \subseteq M$ with $\mathbb{P}(M') = 1$ such that for all $\omega \in M'$, $\sup_{x\in C(\omega)} |f(x) - f^N(x)| \to 0$. Easy to see that x^* is the unique minimum of f on $C(\omega)$ and x^N is a minimum of f^N on $C(\omega)$. Following the same proof of Lemma C.2.2, we know for all $\omega \in M'$, $x^N(\omega) \to x^*(\omega)$, which implies $x^N \xrightarrow{a.s} x^*$. ■

In the remainder of this proof, we aim to show that Γ^N, Γ, $g^N \equiv g^*(s^N)$ and g^* solving $u(g^* - \overline{m}) = u(g^* - \underline{m})$ satisfy the conditions of Lemma C.3.1. We do so in three steps, one for each condition in the lemma. This allows us to obtain that $g^*(s^N) \xrightarrow{a.s} g^*$.

Step 1. $\sup_{g\in C} |\Gamma(g) - \Gamma^N(g)| \xrightarrow{a.s.} 0$, **as** $n \to \infty$, **for all compact sets** $C \subset \mathbb{R}$.

Step 1.1. $\sup_{g\in\mathbb{R}} |\tilde{\Gamma}^N(g) - \Gamma^N(g)| \xrightarrow{a.s.} 0$.

We start by using the auxiliary function $\tilde{\Gamma}^N$. As N grows to infinity, the gap between Γ^N and $\tilde{\Gamma}^N$ shrinks uniformly. We prove this statement next. Start by noticing that, for fixed P^N, for any $p \in \mathbb{P}^a(a^N)$:

$$\Gamma^N(g) \geq \mathbb{E}_p\left[u(g - \theta)\right] \geq u\left(g - \mathbb{E}_p[\theta]\right),$$

where we use convexity of u for the second inequality. Then, by taking max over $p \in \mathbb{P}^a(a^N)$ we obtain $\tilde{\Gamma}^N(g) \leq \Gamma^N(g)$.

Now, for each g, $\tilde{\theta}$ and $q \in [\underline{m}, \overline{m}]$, Taylor's rule implies existence of $\omega(g, \tilde{\theta}, q)$ such that:

$$u(g - \tilde{\theta}) = u\left(g - q\right) + u'(g - \omega(g, \tilde{\theta}, q))\left(\tilde{\theta} - q\right).$$

By the implicit function theorem, $\omega(\cdot, \theta, \cdot)$ is a differentiable, and thus continuous function.

If there exists $p \in \mathbb{P}^{\boldsymbol{a}}(a^N)$ with $\mathbb{E}_p[\theta] = q$, we can take expectations with respect to p in the above equation to obtain:

$$\mathbb{E}_p[u(g-\tilde{\theta})] = u(g-\mathbb{E}_p[\theta]) + \mathbb{E}_p\left[u'(g,\omega(g,\tilde{\theta},\mathbb{E}_p[\theta]))\left(\tilde{\theta}-\mathbb{E}_p[\theta]\right)\right].$$

We can then use subadditivity of the max operator to obtain:

$$\begin{aligned}\Gamma^N(g) &\leq \max_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} u(g-\mathbb{E}_p[\theta]) + \max_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} \mathbb{E}\left[u'\left(g-\omega(g,\tilde{\theta},\mathbb{E}_p[\theta])\right)\left(\tilde{\theta}-\mathbb{E}_p[\theta]\right)\right] \\ &= \tilde{\Gamma}^N(g) + \max_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} \mathbb{E}\left[u'\left(g-\omega(g,\tilde{\theta},\mathbb{E}_p[\theta])\right)\left(\tilde{\theta}-\mathbb{E}_p[\theta]\right)\right]\end{aligned}$$

Fix a compact set C. Define $v(\tilde{\theta}) = \max_{g\in C, q\in[\underline{m},\overline{m}]} u'\left(g-\omega(g,\tilde{\theta},q)\right)$. Which is guaranteed to be well-defined by continuity of u' and ω. We then have:

$$0 \leq \Gamma^N(g) - \tilde{\Gamma}^N(g) \leq \max_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} \mathbb{E}_p\left[v(\tilde{\theta})\left(\tilde{\theta}-\mathbb{E}_p[\theta]\right)\right].$$

Notice that neither bound depends on g within this compact set. On top of that, the upper bound converges to zero. To see that, recall that all signals are informative - $\underline{\rho} > 0$. That implies every $p^N \in \mathbb{P}^{\boldsymbol{a}}(a^N)$ have an almost-sure convergent subsequence to a degenerate distribution. Therefore, $\tilde{\theta} - \mathbb{E}_p[\theta] \to 0$ almost surely. That implies:

$$\sup_{g\in C} |\Gamma^N(g) - \tilde{\Gamma}^N(g)| \xrightarrow{a.s.} 0.$$

Step 1.2. $\sup_{g\in C} |\Gamma(g) - \tilde{\Gamma}^N(g)| \xrightarrow{a.s.} 0$, **as** $n \to \infty$, **for all compact sets** $C \subset \mathbb{R}$.

Recall that we can write $\tilde{\Gamma}^N(g) = \max\{u(g-\overline{m}^N), u(g-\underline{m}^N)\}$. Also by Theorem 3.3.5, $\overline{m}^N \xrightarrow{a.s.} \overline{m}$ and $\underline{m}^N \xrightarrow{a.s.} \underline{m}$. We use the following lemma:

Lemma C.3.2 *Let f^N, g^N, f, g for $N \in \mathbb{N}$ be functions from $D \subset \mathbb{R}$ into the reals, and let $h^N = \max\{f^N, g^N\}$ and $h = \max\{f, g\}$. If $\sup_x |f^N - f| \to 0$ and $\sup_x |g^N - g| \to 0$ then, $\sup_x |h^N - h| \to 0$.*

Proof. For any fixed ε there exist N_f and N_g such that, for all $x \in D$:

$$|f^N(x) - f(x)| < \varepsilon \quad \text{if} \quad N \geq N_f,$$

$$|g^N(x) - g(x)| < \varepsilon \text{ if } N \geq N_g.$$

Take $N \geq \tilde{N} = \max\{N_f, N_g\}$. We then have:

$$\begin{aligned} h(x) &\leq (f^N(x) + \varepsilon)\mathbb{1}_{f(x)\geq g(x)} + (g^N(x) + \varepsilon)\mathbb{1}_{g(x)\geq f(x)} \\ &\leq h^N(x) + \varepsilon. \end{aligned}$$

where the second inequality comes from the definition of h^N. By the same logic, inverting the roles of h and h^N:

$$\begin{aligned} h^N(x) &\leq (f(x) + \varepsilon)\mathbb{1}_{f^N(x)\geq g^N(x)} + (g(x) + \varepsilon)\mathbb{1}_{g^N(x)\geq f^N(x)} \\ &\leq h(x) + \varepsilon. \end{aligned}$$

By joining the two inequalities above: $|h(x) - h^N(x)| \leq \varepsilon$ for all $N \geq \tilde{N}$. Because x is arbitrary, we have our result. ■

In order to apply the result above, notice that $\sup_{g\in C} |u(g - x) - u(g - y)|$ is a continuous function of x and, thus, converges to 0 as $x \to y$. Thus, $\sup_{g\in C} |u(g - \overline{m}^N) - u(g - \overline{m})| \xrightarrow{a.s.} 0$ and similarly $\sup_{g\in C} |u(g - \underline{m}^N) - u(g - \underline{m})| \xrightarrow{a.s.} 0$. Therefore, applying the above lemma, defining $f^N(x) = u(x - \overline{m}^N)$ and $g^N(x) = u(x - \underline{m}^N)$ gives us our result.

Step 1.3. $\sup_{g\in C} |\Gamma(g) - \Gamma^N(g)| \xrightarrow{a.s.} 0$, **as** $n \to \infty$, **for all compact sets** $C \subset \mathbb{R}$.

This is directly implied by the previous two steps.

Step 2. g^* **such that** $u(g^* - \underline{m}) = u(g^* - \overline{m})$ **is the unique minimum of** Γ.

Recall that $\Gamma(g) = \max\{u(g - \underline{m}), u(g - \overline{m})\}$. First, notice that g^* that minimizes Γ must be in $[\underline{m}, \overline{m}]$. Assume, for a contradiction, that $\min \Gamma(g) = u(g^* - \underline{m}) > u(g^* - \overline{m})$. By continuity of u, we can choose $\underline{m} < g' < g^*$ such that $u(g' - \underline{m}) > u(g' - \overline{m})$, that is, $\Gamma(g') = u(g' - \underline{m})$. Because u is strictly convex and minimized at 0, it must be that $u(g^* - \underline{m}) > u(g' - \underline{m})$. But then, $\Gamma(g') < \Gamma(g^*)$, which is a contradiction. A similar contradiction is found if we assume $\min \Gamma(g) = u(g^* - \overline{m}) < u(g^* - \underline{m})$. Thus, the equality must hold.

Step 3. The sequence g^N is uniformly bounded almost everywhere. For an observable realization a^N, recall that $\underline{m}^N = \min_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} \mathbb{E}[\theta]$ and, symmetrically, $\overline{m}^N = \max_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} \mathbb{E}[\theta]$. Assume, for a contradiction, that there is an event M with probability 1, such that g^N is unbounded. If that's the case, up to a subsequence, we have: $g^N > N$. Then, by strict convexity of u we have:

$$\Gamma(g^N) = \max_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} \mathbb{E}_p[u(g^N - \theta)] \geq \max_{p\in\mathbb{P}^{\boldsymbol{a}}(a^N)} u(g^N - \mathbb{E}_p[\theta]) \geq u(g^N - \overline{m}^N).$$

Now, because $\overline{m}^N \xrightarrow{a.s.} \overline{m}$, we can choose an event $M' \subset M$, also with probability one, in which m^N. That implies, with the unboundedness of g^N and strict convexity of u, that the lower bound above diverges, so $\Gamma(g^N)$ is unbounded. To show that g^N cannot be optimal, it suffices to show that there is a sequence x^N such that $\Gamma(x^N)$ is bounded in this event. For any real a, take the sequence $x^N = a$ for all N. Because $\Gamma^N \xrightarrow{a.s.} \Gamma$ uniformly in any compact set, we have that, for a further event $M'' \subset M'$, with probability 1, that for any ε, for sufficiently large N,

$$\Gamma^N(a) < \Gamma(a) + \varepsilon.$$

Thus, $\Gamma^N(a)$ is a bounded sequence, proving that, for sufficiently large N:

$$\Gamma^N(a) < \Gamma^N(g^N),$$

which is the contradiction that we were seeking.

C.4 Proof of Corollary 3.4.3

Define

$$\overline{\zeta}(m) = \frac{\underline{\rho}\int_{-\infty}^{m} x dF(x) + \bar{\rho}\int_{m}^{\infty} x dF(x)}{\underline{\rho}F(m) + \bar{\rho}\,(1 - F(m))}, \quad \underline{\zeta}(m) = \frac{\overline{\rho}\int_{-\infty}^{m} x dF(x) + \underline{\rho}\int_{m}^{\infty} x dF(x)}{\overline{\rho}F(m) + \underline{\rho}\,(1 - F(m))}.$$

Clearly, $\overline{\zeta}(\overline{m}) = \overline{m}$ and $\underline{\zeta}(\underline{m}) = \underline{m}$. Because F is symmetric around θ, for $m \in \mathbb{R}$:

$$\begin{aligned}\overline{\zeta}(2\theta - m) &= \frac{\underline{\rho}\int_{-\infty}^{2\theta-m} x dF(x) + \bar{\rho}\int_{2\theta-m}^{\infty} x dF(x)}{\underline{\rho}F(2\theta - m) + \bar{\rho}\,(1 - F(2\theta - m))} \\ &= 2\theta - \frac{\overline{\rho}\int_{-\infty}^{m} x dF(x) + \underline{\rho}\int_{m}^{\infty} x dF(x)}{\overline{\rho}F(m) + \underline{\rho}\,(1 - F(m))} = 2\theta - \underline{\zeta}(m).\end{aligned}$$

Then, $2\theta - \underline{m} = 2\theta - \underline{\zeta}(\underline{m}) = \overline{\zeta}(2\theta - \underline{m})$. But because $\overline{m}$ is the unique fixed point of $\overline{\zeta}$: $\overline{m} = 2\theta - \underline{m}$, and we are done.

C.5 Proof of Proposition 3.5.1

We go through the monotonicity for $\overline{m}$, a symmetric argument holds for $\underline{m}$. Define $k_\eta(a)$ as

$$k_\eta(a) \equiv \mathbb{E}[\theta](a) = \frac{\underline{\rho}\int_{-\infty}^{a} x f(x)dx + \overline{\rho}\int_{a}^{\infty} x f(x)dx}{\underline{\rho}F(a) + \overline{\rho}\,(1 - F(a))} = \frac{\int_{-\infty}^{a} x f(x)dx + \eta\int_{a}^{\infty} x f(x)dx}{F(a) + \eta\,(1 - F(a))}.$$

For convenience we can rewrite $k_\eta(a)$ as

$$k_\eta(a) = \frac{F(a)\mathbb{E}[x|x < a] + \eta(1 - F(a))\mathbb{E}[x|x \geq a]}{F(a) + \eta(1 - F(a))}.$$

We know that

$$\overline{m} = \arg\max_{a\in\mathbb{R}} k_\eta(a) \qquad \textit{and} \qquad \overline{m} = \max_{a\in\mathbb{R}} k_\eta(a).$$

Then, via the envelope theorem we have

$$\frac{d\overline{m}}{d\eta} = \frac{dk_\eta(\overline{m})}{d\eta} = \frac{F(\overline{m})(1 - F(\overline{m}))}{\left(F(\overline{m}) + \eta(1 - F(\overline{m}))\right)^2}\left(\mathbb{E}[x|x \geq \overline{m}] - \mathbb{E}[x|x < \overline{m}]\right) > 0.$$

Hence, $\overline{m}$ monotonically increases in η. Similarly, $\underline{m}$ monotonically decreases in η.

Now we want to show that as $\eta \to +\infty$, $\overline{m} \to \infty$. First note

$$\lim_{\eta\to\infty} k_\eta(a) = \mathbb{E}[x|x \geq a] > a.$$

The last inequality follows from the full support of the distribution. For any $z \in \mathbb{R}$ we want to show that $\exists\, \tilde{\eta}$ such that $k_{\tilde{\eta}}(\overline{m}) \geq z$. From the above limit, we know that $\exists \tilde{\eta}$ such that $k_{\tilde{\eta}}(z) > z$. Because $\overline{m} = \arg\max_{a \in \mathbb{R}} k_\eta(a)$ we know that $k_{\tilde{\eta}}(\overline{m}) \geq k_{\tilde{\eta}}(z) > z$. Similarly, we can show that as $\eta \to -\infty$, $\underline{m} \to -\infty$.

Finally, we argue that as $\eta \to 1$, $\overline{m} - \underline{m} \to 0$. When $\eta \to 1$, $k_\eta(a)$ reduces to the unconditional expected value for any a. Similarly, the optimization problem that determines $\underline{m}$ reduces to the unconditional expected value, completely unaffected by a. Thus, as $\eta \to 1$ both $\overline{m}$ and $\underline{m}$ converge to the unconditional expectation.

C.6 Proof of Proposition 3.5.3

We use k_η as define in the Proof of Proposition 3.5.1. Fix any $\eta > 1$ and $C > 0$ and consider the family of distributions over precisions $G_\rho = \delta_\rho$, $\rho > 0$. Given the state, the real signals are distributes according to $s \sim \mathcal{N}(\theta, \frac{1}{\rho})$. Let F_ρ be the distribution of s.

We use k_η as defined in the Proof of Proposition 3.5.1, but we call it k_ρ here, as we are interested in precisions instead. We can then use the formula for conditional means of normal distributions to write:

$$k_\rho(a) = \theta + \frac{\frac{1}{\rho}(\eta - 1) f_\rho(a)}{F_\rho(a) + \eta(1 - F_\rho(a))}.$$

For each a, as ρ converges to 0, the denominator converges to $\frac{1+\eta}{2}$, while the numerator converges to infinity. Because $\overline{m}_\rho = \max_a k_\rho(a)$, there is a sufficiently low ρ so that $\overline{m}_\rho$ is arbitrarily large. A similar argument can be made for $\underline{m}_\rho$. Now, by Example 3.4.4, we have

$$g^* = \frac{\underline{m}_\rho + \lambda \overline{m}_\rho}{1 + \sqrt{\lambda}}.$$

Thus, $|g^* - \theta| = \frac{|1-\sqrt{\lambda}|(\overline{m}_\rho - \underline{m}_\rho)}{1+\sqrt{\lambda}}$ that can be made arbitrarily large for ρ sufficiently close to zero, which concludes our proof. Note that, although this result was proved using a degenerate G, it can easily be extended to continuous G's.

C.7 Proof of Proposition 3.5.4

We go through the proof for $\overline{m}$, a symmetric argument holds for $\underline{m}$. From Theorem 3.3.4 we know that

$$\overline{m} = \frac{\underline{\rho}\int_{-\infty}^{\overline{m}} \boldsymbol{s^a}(x,\underline{\rho})dF(x) + \overline{\rho}\int_{\overline{m}}^{\infty} \boldsymbol{s^a}(x,\overline{\rho})dF(x)}{\underline{\rho}F(\overline{m}) + \overline{\rho}\,(1-F(\overline{m}))}$$

where $\boldsymbol{s^a}(x,\rho)$ is the inverted signal given action x and conjectured precision ρ, and $F(x)$ is the distribution of the observables. In the observable actions case the inverted signal is simply $\boldsymbol{s^a}\,(a_i,\hat{\rho}_i) = a_i + \frac{\rho_\mu}{\hat{\rho}_i}(a_i-\mu)$, thus the above equation becomes

$$\begin{aligned}\overline{m}_a &= \frac{\underline{\rho}\int_{-\infty}^{\overline{m}_a}\left(x+\frac{\rho_\mu}{\underline{\rho}}(x-\mu)\right)dF(x) + \overline{\rho}\int_{\overline{m}_a}^{\infty}\left(x+\frac{\rho_\mu}{\overline{\rho}}(x-\mu)\right)dF(x)}{\underline{\rho}F(\overline{m}_a) + \overline{\rho}\,(1-F(\overline{m}_a))} \\ &= \frac{\underline{\rho}\int_{-\infty}^{\overline{m}_a} xdF(x) + \overline{\rho}\int_{\overline{m}_a}^{\infty} xdF(x) + \int_{-\infty}^{\infty}\rho_\mu(x-\mu)dF(x)}{\underline{\rho}F(\overline{m}_a) + \overline{\rho}\,(1-F(\overline{m}_a))}.\end{aligned}$$

Recall that actions are $\boldsymbol{a}(s_i,\rho_i) = \frac{\rho_\mu\mu+\rho_i s_i}{\rho_\mu+\rho_i}$, where ρ_i is the true not conjectured precision of the agent. Thus, given ρ_i the expected value of the observable is $\frac{\rho_\mu\mu+\rho_i\theta}{\rho_\mu+\rho_i}$, since the signals normally distributed around θ. Also,

$$F(x) = \int_{[\underline{\rho},\overline{\rho}]} F_\rho\Big(\boldsymbol{s^a}(x,\rho)\Big)dG(\rho).$$

This leads to

$$\int_{-\infty}^{\infty}\rho_\mu(x-\mu)dF(x) = (\theta-\mu)\int\frac{\rho\rho_\mu}{\rho_\mu+\rho}dG(\rho) := c.$$

C.8 Proof of Proposition 3.5.7

Recall from Proposition 3.5.4 that the bounds of the limiting posterior set are given by

$$\overline{m}_a = \frac{\underline{\rho}\int_{-\infty}^{\overline{m}_a} x dF(x) + \bar{\rho}\int_{\overline{m}_a}^{\infty} x dF(x) + c}{\underline{\rho}F(\overline{m}_a) + \bar{\rho}\left(1 - F(\overline{m}_a)\right)}, \qquad \underline{m}_a = \frac{\overline{\rho}\int_{-\infty}^{\underline{m}_a} x dF(x) + \underline{\rho}\int_{\underline{m}_a}^{\infty} x dF(x) + c}{\overline{\rho}F(\underline{m}_a) + \underline{\rho}\left(1 - F(\underline{m}_a)\right)} \tag{C.4}$$

where $c = \frac{\rho\rho_\mu}{\rho_\mu+\rho}(\theta - \mu)$.

The optimal guess is $m_a = \frac{\overline{m}_a + \underline{m}_a}{2}$. When $\theta = \mu$, $c = 0$ and by Corollary 3.4.3 $m_a = \theta = \mu$ and the observer guesses correctly. From now on, we first focus on the case where $\theta > \mu$. The case for $\theta < \mu$ is symmetric.

Denote $\overline{G}(z) = \underline{\rho}F(z) + \bar{\rho}\left(1 - F(z)\right)$ and $\underline{G}(z) = \overline{\rho}F(z) + \underline{\rho}\left(1 - F(z)\right)$. Rearranging the first equation and using integration by parts, we get

$$\begin{aligned}
\overline{m}_a\overline{G}(\overline{m}_a) &= \underline{\rho}\left(xF(x)\Big|_{-\infty}^{\overline{m}_a} - \int_{-\infty}^{\overline{m}_a} F(x)dx\right) + \overline{\rho}\left(-x\left(1 - F(x)\right)\Big|_{\overline{m}_a}^{\infty} + \int_{\overline{m}_a}^{\infty}\left(1 - F(x)\right)dx\right) + c \\
&= \underline{\rho}\left(\overline{m}_a F(\overline{m}_a) - \int_{-\infty}^{\overline{m}_a} F(x)dx\right) + \overline{\rho}\left(\overline{m}_a\left(1 - F(\overline{m}_a)\right) + \int_{\overline{m}_a}^{\infty}\left(1 - F(x)\right)dx\right) + c \\
&= \overline{m}_a\overline{G}(\overline{m}_a) - \left(\underline{\rho}\int_{-\infty}^{\overline{m}_a} F(x)dx - \overline{\rho}\int_{\overline{m}_a}^{\infty}\left(1 - F(x)\right)dx\right) + c.
\end{aligned}$$

This implies

$$\underline{\rho}\int_{-\infty}^{\overline{m}_a} F(x)dx - \overline{\rho}\int_{\overline{m}_a}^{\infty}\left(1 - F(x)\right)dx = c. \tag{C.5}$$

A symmetric argument for $\underline{m}_a$ shows that

$$\overline{\rho}\int_{-\infty}^{\underline{m}_a} F(x)dx - \underline{\rho}\int_{\underline{m}_a}^{\infty}\left(1 - F(x)\right)dx = c. \tag{C.6}$$

Taking the derivative with respect to the state θ on both sides of equation C.5 and equation C.6, we get

$$\frac{d\,\overline{m}_a}{d\,\theta} = \frac{\rho}{\rho_\mu + \rho} + \frac{\rho_\mu}{\rho_\mu + \rho}\frac{\rho}{\overline{G}(\overline{m}_a)} \qquad \frac{d\,\underline{m}_a}{d\,\theta} = \frac{\rho}{\rho_\mu + \rho} + \frac{\rho_\mu}{\rho_\mu + \rho}\frac{\rho}{\underline{G}(\underline{m}_a)}$$

The derivative of the optimal guess $m_a = \frac{\overline{m}_a + \underline{m}_a}{2}$ with respect to θ is then:

$$\frac{dm_a}{d\theta} = \frac{\rho}{\rho_\mu + \rho} + \frac{\rho_\mu}{\rho_\mu + \rho}\frac{\rho}{2}\left(\frac{1}{\overline{G}(\overline{m}_a)} + \frac{1}{\underline{G}(\underline{m}_a)}\right) \tag{C.7}$$

Recall that H is normally distributed and denote its density function as h. Then, we can use the derivative of the optimal bounds obtained above to calculate:

$$\begin{aligned}
\frac{dF(\overline{m}_a)}{d\theta} &= \frac{\partial F(\overline{m}_a)}{\partial \overline{m}_a}\frac{d\overline{m}_a}{d\theta} + \frac{\partial F(\overline{m}_a)}{\partial \theta} = -\frac{\rho_\mu}{\rho_\mu + \rho}\frac{\rho}{\overline{G}(\overline{m}_a)} f(\overline{m}_a), \\
\frac{dF(\underline{m}_a)}{d\theta} &= \frac{\partial F(\underline{m}_a)}{\partial \underline{m}_a}\frac{d\underline{m}_a}{d\theta} + \frac{\partial F(\underline{m}_a)}{\partial \theta} = -\frac{\rho_\mu}{\rho_\mu + \rho}\frac{\rho}{\underline{G}(\underline{m}_a)} f(\underline{m}_a).
\end{aligned}$$

It then follows that

$$\begin{aligned}
\frac{d^2 m_a}{d\theta^2} &= \frac{\overline{\rho} - \underline{\rho}}{2}\left(\frac{\rho\rho_\mu}{\rho_\mu + \rho}\right)^2 \left(\frac{f(\overline{m}_a)}{\overline{G}^3(\overline{m}_a)} - \frac{f(\underline{m}_a)}{\underline{G}^3(\underline{m}_a)}\right) \\
&= \frac{\overline{\rho} - \underline{\rho}}{2}\left(\frac{\rho\rho_\mu}{\rho_\mu + \rho}\right)^2 \left(\left(\frac{f(\overline{m}_a)}{\overline{G}(\overline{m}_a)} - \frac{f(\underline{m}_a)}{\underline{G}(\underline{m}_a)}\right)\frac{1}{\underline{G}^2(\underline{m}_a)} + \frac{f(\overline{m}_a)}{\overline{G}(\overline{m}_a)}\left(\frac{1}{\overline{G}^2(\overline{m}_a)} - \frac{1}{\underline{G}^2(\underline{m}_a)}\right)\right).
\end{aligned}$$

Lemma C.8.1 $\left(\frac{1}{\overline{G}^2(\overline{m}_a)} - \frac{1}{\underline{G}^2(\underline{m}_a)}\right) > 0$ *whenever* $\theta > \mu$

Proof of Lemma C.8.1. The statement is equivalent to $\underline{G}(\underline{m}_a) > \overline{G}(\overline{m}_a)$, which is also equivalent to $F(\overline{m}_a) + F(\underline{m}_a) > 1$. Since H is symmetric around $\frac{\rho\theta + \rho_\mu\mu}{\rho + \rho_\mu}$, the latter is true if and only if $m_a > \frac{\rho\theta + \rho_\mu\mu}{\rho + \rho_\mu}$. We show that this is the case. Define

$$\begin{aligned}
\overline{\zeta}(z,u) &= \frac{\underline{\rho}\int_{-\infty}^{z} x dF(x) + \bar{\rho}\int_{z}^{\infty} x dF(x) + u}{\underline{\rho}F(z) + \bar{\rho}\,(1 - F(z))}, \\
\underline{\zeta}(z,u) &= \frac{\overline{\rho}\int_{-\infty}^{z} x dF(x) + \underline{\rho}\int_{z}^{\infty} x dF(x) + u}{\overline{\rho}F(z) + \underline{\rho}\,(1 - F(z))}.
\end{aligned}$$

We know $\overline{m}_a = \overline{\zeta}(\overline{m}_a, c)$, and it was previously proved that $\overline{m}_a$ maximizes $\overline{\zeta}(\overline{m}_a, c)$. By the envelope theorem we have:

$$\frac{d\overline{\zeta}(\overline{m}_a, c)}{du} = \frac{\partial\overline{\zeta}(\overline{m}_a, c)}{\partial u} = \frac{1}{\underline{\rho}F(\overline{m}_a) + \bar{\rho}\,(1 - F(\overline{m}_a))} > 0.$$

A similar argument implies that $\frac{\underline{\zeta}(\underline{m}_a, u)}{du} > 0$, for all $u \in \mathbb{R}$. Finally, by an equivalent argument to the proof of Corollary 3.4.3, we have $\frac{\overline{\zeta}(\overline{m}_a, 0) + \underline{\zeta}(\underline{m}_a, 0)}{2} = \int x dH = \frac{\rho\theta + \rho_\mu \mu}{\rho + \rho_m u}$. Then, if $\theta > \mu$ - which implies $c > 0$:

$$m_a = \frac{\overline{m}_a + \underline{m}_a}{2} = \frac{\overline{\zeta}(\overline{m}_a, c) + \underline{\zeta}(\underline{m}_a, c)}{2} > \frac{\overline{\zeta}(\overline{m}_a, 0) + \underline{\zeta}(\underline{m}_a, 0)}{2}.$$

This concludes the proof of the lemma. ■

Therefore,

$$\left(\frac{f(\overline{m}_a)}{\overline{G}(\overline{m}_a)} - \frac{f(\underline{m}_a)}{\underline{G}(\underline{m}_a)}\right) \geq 0 \implies \frac{d^2 m_a}{d\theta^2} > 0. \tag{C.8}$$

We next consider the partial derivative of the optimal guess with respect to ρ. We start with an alternative implicit function of $\overline{m}_a$ and $\underline{m}_a$. Notice that if f as the density function of a normal distribution with mean $\tilde{\mu}$ and variance $\tilde{\sigma}^2$, then $\frac{\partial f(x)}{\partial x} = -\frac{x - \tilde{\mu}}{\tilde{\sigma}^2} f(x)$. This implies $x f(x) = \tilde{\mu} f(x) - \tilde{\sigma}^2 \frac{\partial f(x)}{\partial x}$. Plugging this into the initial implicit functions C.4, we get

$$\begin{aligned}
\overline{m}_a &= \frac{\rho_\mu \mu + \rho\theta}{\rho_\mu + \rho} + \frac{c}{\overline{G}(\overline{m}_a)} + (\overline{\rho} - \underline{\rho})\frac{\rho}{(\rho_\mu + \rho)^2}\frac{f(\overline{m}_a)}{\overline{G}(\overline{m}_a)}, \\
\underline{m}_a &= \frac{\rho_\mu \mu + \rho\theta}{\rho_\mu + \rho} + \frac{c}{\underline{G}(\underline{m}_a)} - (\overline{\rho} - \underline{\rho})\frac{\rho}{(\rho_\mu + \rho)^2}\frac{f(\underline{m}_a)}{\underline{G}(\underline{m}_a)}.
\end{aligned}$$

By definition of m_a, we have

$$m_a = \theta + (\theta - \mu)\left(\frac{dm_a}{d\theta} - 1\right) + \frac{(\overline{\rho} - \underline{\rho})\rho}{2(\rho_\mu + \rho)^2}\left(\frac{f(\overline{m}_a)}{\overline{G}(\overline{m}_a)} - \frac{f(\underline{m}_a)}{\underline{G}(\underline{m}_a)}\right). \tag{C.9}$$

Based on the implicit function theorem, we can calculate the following derivative:

$$\frac{dm_a}{d\rho} = \frac{\rho_\mu(m_a - \mu) + \rho(\theta - m_a)}{2\rho^2 + 2\rho_\mu\rho} + \frac{c}{2}\frac{\rho_\mu + (\rho_\mu + \rho)\rho}{(\rho_\mu + \rho)\rho}\left(\frac{1}{\overline{G}(\overline{m}_a)} + \frac{1}{\underline{G}(\underline{m}_a)}\right).$$

As $\theta > \mu$, it is easy to show that $m_a > \mu$ and $c > 0$. This leads to the following result.

$$\theta > \mu \quad \text{and} \quad m_a \leq \theta \quad \Longrightarrow \quad \frac{dm_a}{d\rho} > 0. \tag{C.10}$$

Note that the last term of $\frac{dm_a}{d\rho}$, $\left(\frac{1}{\overline{G}(\overline{m}_a)} + \frac{1}{\underline{G}(\underline{m}_a)}\right)$ can be rewritten as $\left(\frac{dm_a}{d\theta} - \frac{\rho}{\rho_\mu+\rho}\right)\frac{\rho_\mu+\rho}{\rho_\mu}\frac{2}{\rho}$. Let $\kappa_1 = \frac{1}{2\rho^2+2\rho_\mu\rho}$ and $\kappa_2 = \frac{\rho_\mu+(\rho_\mu+\rho)\rho}{\rho_\mu\rho^2}$, then:

$$\frac{d^2m_a}{d\rho\, d\theta} = \rho_\mu\kappa_1\frac{dm_a}{d\theta} - \rho\kappa_1\left(\frac{dm_a}{d\theta} - 1\right) + \frac{\rho\rho_\mu}{\rho_\mu+\rho}\kappa_2\left(\frac{dm_a}{d\theta} - \frac{\rho}{\rho_\mu+\rho}\right) + c\kappa_2\frac{d^2m_a}{d\theta^2} \tag{C.11}$$

We know that $\frac{dm_a}{d\theta} > \frac{\rho}{\rho_\mu+\rho} > 0$ and when $\theta = \mu$, $\frac{d^2m_a}{d\theta^2} = 0$. This leads to the following result:

$$\theta = \mu \quad \text{and} \quad \frac{dm_a}{d\theta} \leq 1 \quad \Longrightarrow \quad \frac{d^2m_a}{d\rho\, d\theta} > 0. \tag{C.12}$$

To make it clear that the optimal guess depends on θ and ρ, we sometimes denote $\underline{m}_a$, $\overline{m}_a$ and m_a as $\underline{m}_a(\rho,\theta)$, $\overline{m}_a(\rho,\theta)$ and $m_a(\rho,\theta)$. Notice that $\tilde{\rho}$ is determined by forcing $\frac{dm_a}{d\theta}$ to approach 1 when θ goes to infinity, while at $\tilde{\tilde{\rho}}$ we have $\frac{dm_a}{d\theta}(\tilde{\tilde{\rho}},\mu) = 1$.

The rest of the proof will be divided by the following lemmas. We will fix μ and consider the case with $\theta \geq \mu$.

Lemma C.8.2 *For any given ρ, if $m_a(\rho,\hat{\theta}) > \hat{\theta}$ and $\frac{dm_a}{d\theta}(\rho,\hat{\theta}) > 1$, then $m_a(\rho,\theta) > \theta$ for all $\theta > \hat{\theta}$.*

Proof of Lemma C.8.2. Fix ρ. Assume that there exists $\hat{\theta}$, $m_a(\rho,\hat{\theta}) > \hat{\theta}$ and $\frac{dm_a}{d\theta}(\rho,\hat{\theta}) > 1$. Suppose by contradiction that there exists some $\overline{\theta} > \hat{\theta}$ such that

$m_a(\rho, \overline{\theta}) = \overline{\theta}$. By continuity of $\frac{dm_a}{d\theta}$, there exists $\theta' < \theta'' \in (\hat{\theta}, \overline{\theta}]$ where $\frac{dm_a}{d\theta}(\rho, \theta') = 1$ and $\frac{dm_a}{d\theta}(\rho, \theta'') < 1$. By continuity of m_a, $m_a(\rho, \theta') > \theta'$.

At θ', equation (C.9) implies $\left(\frac{f(\overline{m}_a)}{\overline{G}(\overline{m}_a)} - \frac{f(\underline{m}_a)}{\underline{G}(\underline{m}_a)}\right) > 0$, which guarantees $\frac{d^2 m_a}{d\theta^2}(\rho, \theta') > 0$. This implies that for a neighborhood to the right of θ', $\frac{dm_a}{d\theta} > 1$. Notice that this holds for any $\theta \in [\hat{\theta}, \overline{\theta}]$ with $\frac{dm_a}{d\theta}(\rho, \theta) = 1$. Thus $\frac{dm_a}{d\theta}(\rho, \theta) \geq 1$ for all $\theta \in [\hat{\theta}, \overline{\theta}]$, which contradicts the assumption that $m_a(\rho, \overline{\theta}) = \overline{\theta}$. As a result, we know $m_a(\rho, \theta) > \theta$ for $\theta > \hat{\theta}$. This concludes the proof of the lemma. ■

Lemma C.8.3 *For any given ρ, if there exists $\theta^* > \mu$ such that $m_a(\rho, \theta^*) = \theta^*$ and $m_a(\rho, \theta) < \theta$ for all $\mu < \theta < \theta^*$, then $m_a(\rho, \theta) > \theta$ for $\theta > \theta^*$.*

Proof of Lemma C.8.3. Suppose there exists $\theta^* > \mu$ such that $m_a(\rho, \theta^*) = \theta^*$ and $m_a(\rho, \theta) < \theta$ for $\mu < \theta < \theta^*$. This implies $\frac{dm_a}{d\theta}(\rho, \theta^*) \geq 1$. Again by equation (C.9), we know $\left(\frac{f(\overline{m}_a)}{\overline{G}(\overline{m})} - \frac{f(\underline{m}_a)}{\underline{G}(\underline{m}_a)}\right) > 0$, which leads to $\frac{d^2 m_a}{d\theta^2}(\rho, \theta^*) > 0$ by (C.8). Then for any θ in a small neighborhood to the right of θ^*, $\frac{dm_a}{d\theta}(\rho, \theta) > 1$ and $m_a(\rho, \theta) > \theta$. By Lemma C.8.2. This concludes the proof of the lemma and the proposition. ■

References

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2018): "Multiattribute Utility Theory, Intertemporal Utility, and Correlation Aversion," *International Economic Review*, 59, 537–555.

Attema, A. E., O. L'haridon, and G. van de Kuilen (2019): "Measuring Multivariate Risk Preferences in the Health Domain," *Journal of Health Economics*, 64, 15–24.

Bansal, R. and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *The Journal of Finance*, 59, 1481–1509.

Barberis, N., M. Huang, and R. H. Thaler (2006): "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing," *American Economic Review*, 96, 1069–1090.

Barro, R. J. (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, 243–64.

Bergemann, D. and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

Berk, R. H. (1966): "Limiting Behavior of Posterior Distributions when the Model is Incorrect," *The Annals of Mathematical Statistics*, 51–58.

Bikhchandani, S., S. Chatterji, R. Lavi, A. Mu'alem, N. Nisan, and A. Sen (2006): "Weak Monotonicity Characterizes Deterministic Dominant-Strategy Implementation," *Econometrica*, 74, 1109–1132.

Blum, J. R. (1955): "On the Convergence of Empiric Distribution Functions," *The Annals of Mathematical Statistics*.

Bommier, A. (2007): "Risk Aversion, Intertemporal Elasticity of Substitution and Correlation Aversion," *Economics Bulletin*, 4, 1–8.

Bommier, A., A. Chassagnon, and F. Le Grand (2012): "Comparative Risk Aversion: A Formal Approach with Applications to Saving Behavior," *Journal of Economic Theory*, 147, 1614–1641.

Bommier, A., A. Kochov, and F. Le Grand (2017): "On Monotone Recursive Preferences," *Econometrica*, 85, 1433–1466.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012): "Salience Theory of Choice under Risk," *The Quarterly Journal of Economics*, 127, 1243–1285.

Borgers, T. (2015): *An Introduction to the Theory of Mechanism Design*, Oxford University Press.

Bose, S. and A. Daripab (2009): "A Dynamic Mechanism and Surplus Extraction under Ambiguity," *Journal of Economic Theory*, 144, 2084–2114.

Bose, S., E. Ozdenoren, and A. Pape (2006): "Optimal Auctions with Ambiguity," *Theoretical Economics*, 1, 411–438.

Bose, S. and L. Renou (2014): "Mechanism Design With Ambiguous Communication Devices," *Econometrica*, 82, 1853–1872.

Brooks, B. A. and S. Du (2020): "Optimal Auction Design with Common Values: An Informationally-Robust Approach," *Econometrica*, Forthcoming.

Camara, M. (2021): "Computationally Tractable Choice," Working Paper.

Carroll, G. (2015): "Robustness and Linear Contracts," *American Economic Review*, 105, 536–563.

——— (2017): "Robustness and Separation in Multidimensional Screening," *Econometrica*, 85, 453–488.

Castro, L. D. and N. C. Yannelis (2018): "Uncertainty, Efficiency and Incentive Compatibility: Ambiguity Solves the Conflict between Efficiency and Incentive Compatibility," *Journal of Economic Theory*, 177, 678 – 707.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011): "Uncertainty Averse Preferences," *Journal of Economic Theory*, 146, 1275–1330.

——— (2013): "Ambiguity and Robust Statistics," *Journal of Economic Theory*, 148, 974–1049.

Chen, Y.-C. and J. Li (2018): "Revisiting the Foundations of Dominant-Strategy Mechanisms," *Journal of Economic Theory*, 178, 294–317.

Chew, S. H. and L. G. Epstein (1991): "Recursive Utility under Uncertainty," in *Equilibrium Theory in Infinite Dimensional Spaces*, Springer, 352–369.

Chew, S. H., B. Miao, and S. Zhong (2017): "Partial Ambiguity," *Econometrica*, 85, 1239–1260.

Chung, K.-S. and J. Ely (2007): "Foundations of Dominant-Strategy Mechanisms," *Review of Economic Studies*, 74, 447–476.

Condie, S. and J. Ganguli (2017): "The Pricing Effects of Ambiguous Private Information," *Journal of Economic Theory*, 172, 512–557.

DeHardt, J. (1971): "Generalizations of the Glivenko-Cantelli Theorem," *The Annals of Mathematical Statistics*, 42, 2050–2055.

Di Tillio, A., N. Kos, and M. Messner (2017): "The Design of Ambiguous Mechanisms," *The Review of Economic Studies*, 84, 237–276.

Dillenberger, D., D. Gottlieb, and P. Ortoleva (2020): "Stochastic Impatience and the separation of Time and Risk Preferences," Working Paper.

Drèze, J. H. and F. Modigliani (1972): "Consumption Decisions under Uncertainty," *Journal of Economic Theory*, 5, 308–335.

Du, S. (2018): "Robust Mechanisms under Common Valuation," *Econometrica*, 86, 1569–1588.

Ebert, S. and G. van de Kuilen (2016): "Measuring Multivariate Risk Preferences," Working Paper.

Ellis, A. (2021): "Correlation Concern," Working Paper.

Ellis, A. and D. J. Freeman (2020): "Revealing Choice Bracketing," Working Paper.

Ellis, A. and M. Piccione (2017): "Correlation Misperception in Choice," *American Economic Review*, 107, 1264–92.

Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, 70, 643–669.

Enke, B. and F. Zimmermann (2019): "Correlation Neglect in Belief Formation," *The Review of Economic Studies*, 86, 313–332.

Epstein, L. and S. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–69.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014): "How Much Would You Pay to Resolve Long-Run Risk?" *American Economic Review*, 104, 2680–97.

Epstein, L. G. and M. Schneider (2007): "Learning under Ambiguity," *The Review of Economic Studies*, 74, 1275–1303.

——— (2008a): "Ambiguity, Information Quality, and Asset Pricing," *The Journal of Finance*, 63, 197–228.

——— (2008b): "Ambiguity, Information Quality, and Asset Pricing," *The Journal of Finance*, 63, 197–228.

Epstein, L. G. and S. E. Zin (1991): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy*, 99, 263–286.

Esponda, I., D. Pouzo, and Y. Yamamoto (2019): "Asymptotic Behavior of Bayesian Learners with Misspecified Models," *Working Paper*.

Eyster, E. and G. Weizsäcker (2016): "Correlation Neglect in Portfolio Choice: Lab Evidence," Working Paper.

Fishburn, P. C. (1982): *The Foundations of Expected Utility.*, Theory & Decision Library.

Frick, M., R. Iijima, and Y. Ishii (2020): "Stability and Robustness in Misspecified Learning Models," *Working Paper*.

Fudenberg, D., G. Lanzani, and P. Strack (2020): "Limits Points of Endogenous Misspecified Learning," *Working Paper*.

Fudenberg, D., G. Romanyuk, and P. Strack (2017): "Active Learning with a Misspecified Prior," *Theoretical Economics*, 12, 1155–1189.

Giacomini, R. and T. Kitagawa (2021): "Robust Bayesian inference for set-identified models," *Econometrica*, 89, 1519–1556.

Giacomini, R., T. Kitagawa, and H. Uhlig (2019): "Estimation under Ambiguity," *Working Paper*.

Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

——— (1993): "Updating Ambiguous Beliefs," *Journal of Economic Theory*, 59, 33–49.

Gollier, C. (2011): "Portfolio Choices and Asset Prices: The Comparative Statics of Ambiguity Aversion," *The Review of Economic Studies*, 78, 1329–1344.

Gui, H., R. Müller, and R. V. Vohra (2004): "Dominant Strategy Mechanisms with Multidimensional Types," Working Paper.

Guo, H. (2019): "Mechanism Design with Ambiguous Transfers: An Analysis in Finite Dimensional Naive Type Spaces," *Journal of Economic Theory*, 183, 76–105.

Halevy, Y. (2007): "Ellsberg Revisited: An Experimental Study," *Econometrica*, 75, 503–536.

Heidhues, P., B. Koszegi, and P. Strack (2019): "Convergence in Misspecified Learning Models with Endogenous Actions," *Working Paper*.

Huber, P. J. (2004): *Robust statistics*, vol. 523, John Wiley & Sons.

Illeditsch, P. K. (2011): "Ambiguous information, portfolio inertia, and excess volatility," *The Journal of Finance*, 66, 2213–2247.

Kallir, I. and D. Sonsino (2009): "The Neglect of Correlation in Allocation Decisions," *Southern Economic Journal*, 1045–1066.

Ke, S. and Q. Zhang (2020): "Randomization and Ambiguity Aversion," *Econometrica*, 88, 1159–1195.

Keynes, J. M. (1921): *A Treatise on Probability*, Macmillan, London.

Kihlstrom, R. E. and L. J. Mirman (1974): "Risk Aversion with Many Commodities," *Journal of Economic Theory*, 8, 361–388.

Kimball, M. and P. Weil (2009): "Precautionary Saving and Consumption Smoothing across Time and Possibilities," *Journal of Money, Credit and Banking*, 41, 245–284.

Kimball, M. S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica*, 58, 53–73.

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

Knight, F. H. (1921): *Risk, Uncertainty and Profit*, Houghton Mifflin, Boston.

Kreps, D. M. and E. L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 185–200.

Kubler, F., L. Selden, and X. Wei (2020a): "Incomplete Market Demand Tests for Kreps-Porteus-Selden Preferences," *Journal of Economic Theory*, 185, 104973.

——— (2020b): "Time Consistency, Temporal Resolution Indifference and the Separation of Time and Risk," Working Paper.

Lanier, J., B. Miao, J. K.-H. Quah, and S. Zhong (2020): "Intertemporal Consumption with Risk: A Revealed Preference Analysis," Working Paper.

Lanzani, G. (2021): "Correlation Made Simple: Applications to Salience and Regret Theory," *The Quarterly Journal of Economics*.

L.Bodoh-Creed, A. (2012): "Ambiguous Beliefs and Mechanism Design," *Games and Economic Behavior*, 75, 518–537.

Lian, C. (2020): "A Theory of Narrow Thinking," *The Review of Economic Studies*.

Lopomo, G., L. Rigotti, and C. Shannon (2020): "Uncertainty in Mechanism Design," Working Paper.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006): "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," *Econometrica*, 74, 1447–1498.

Marinacci, M. (2002): "Learning from Ambiguous Urns," *Statistical Papers*, 43, 143.

Meissner, T. and P. Pfeiffer (2022): "Measuring Preferences over the Temporal Resolution of Consumption Uncertainty," *Journal of Economic Theory*, 200, 105379.

Mu, X., L. Pomatto, P. Strack, and O. Tamuz (2021): "Background Risk and Small-Stakes Risk Aversion," Working Paper.

Nyarko, Y. (1991): "Learning in Mis-specified Models and the Possibility of Cycles," *Journal of Economic Theory*, 55, 416–427.

Pires, C. P. (2002): "A Rule for Updating Ambiguous Beliefs," *Theory and Decision*, 53, 137–152.

Rabin, M. (2000): "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68, 1281–1292.

Rabin, M. and G. Weizsäcker (2009): "Narrow Bracketing and Dominated Choices," *American Economic Review*, 99, 1508–43.

Read, D., G. Loewenstein, and M. Rabin (1999): "Choice Bracketing," *Journal of Risk and Uncertainty*, 19, 171–197.

Rees-Jones, A., R. Shorrer, and C. J. Tergiman (2020): "Correlation Neglect in Student-to-School Matching," Working Paper.

Roberts, K. (1979): "The Characterization of Implementable Choice Rules," in *Aggregation and Revelation of Preferences*, ed. by J.-J. Laffont, North Holland Publishing Company.

Rochet, J.-C. (1987): "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context," *Journal of Mathematical Economics*, 16, 191–200.

Rockafellar, R. T. (1970): *Convex Analysis*, Princeton University Press, Princeton.

Saito, K. (2015): "Preferences for Flexibility and Randomization under Uncertainty," *American Economic Review*, 105, 1246–1271.

Saks, M. and L. Yu (2005): "Weak Monotonicity Suffices for Truthfulness on Convex Domains," in *Proceedings of the 6th ACM Conference on Electronic Commerce*, ACM, EC '05, 286–293.

Selden, L. (1978): "A New Representation of Preferences over "Certain x Uncertain" Consumption Pairs: The "Ordinal Certainty Equivalent" Hypothesis," *Econometrica*, 1045–1060.

Selden, L. and I. Stux (1978): "Consumption Trees, OCE Utility and the Consumption/Savings Decision," Working Paper.

Shalizi, C. R. (2009): "Dynamics of Bayesian Updating with Dependent Data and Misspecified Models," *Electronic Journal of Statistics*, 3, 1039–1074.

Song, Y. (2018): "Efficient Implementation with Interdependent Valuations and Maxmin Agents," *Journal of Economic Theory*, 176, 693–726.

Thaler, R. (1985): "Mental Accounting and Consumer Choice," *Marketing Science*, 4, 199–214.

Tversky, A. and D. Kahneman (1981): "The Framing of Decisions and the Psychology of Choice," *Science*, 211, 453–458.

Vorjohann, P. (2021): "Reference-Dependent Choice Bracketing," Working Paper.

Weil, P. (1990): "Nonexpected Utility in Macroeconomics," *The Quarterly Journal of Economics*, 105, 29–42.

Wolitzky, A. (2016): "Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade," *Theoretical Economics*, 11, 971–1004.