

Behavioral Network Economics

Soham Phade



Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2021-216

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-216.html>

September 14, 2021

Copyright © 2021, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

I would like to thank my advisor Venkat Anantharam and my other two committee members, Jean Walrand and David Aldous.

Behavioral Network Economics

by

Soham Rajesh Phade

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering — Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Venkat Anantharam, Chair

Professor Jean Walrand

Professor David Aldous

Summer 2021

Behavioral Network Economics

Copyright 2021
by
Soham Rajesh Phade

Abstract

Behavioral Network Economics

by

Soham Rajesh Phade

Doctor of Philosophy in Engineering — Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Venkat Anantharam, Chair

Game theoretic models are prevalent in the study of interactions between autonomous agents. Given the pervasive role of humans as agents in networks (e.g. social networks) and markets (e.g. labor markets), building mechanisms based on presumably more accurate models of human behavior is of great interest both for increasing human welfare and for building more efficient commercial systems that interact with humans. Cumulative prospect theory (CPT), one of the leading models for decision-making under risk and uncertainty, introduced by Kahneman and Tversky, combines several psychological insights into decision theory. Theoretical economics has primarily focused on expected utility theory (EUT) to model human behavior. On the other hand, CPT has been observed to be a better fit in empirical studies, it is a generalization of EUT, and has a nice mathematical formulation convenient for theoretical studies. It provides a way to incorporate psychological aspects into the concrete frameworks of game theory and economics which is required in building large scale systems that are better aligned with human preferences and needs and are also robust to their emotional traits. A systematic and principled approach is needed. This thesis aims to build work in this direction by studying the following three problems through the lens of CPT:

1. resource allocation over networks,
2. notions of equilibrium in non-cooperative games, and
3. mechanism design.

In this thesis, we develop theoretical tools and establish fundamental results that would support real-world applications and future research in behavioral network economics.

To *Mummy* and *Dada*.

Contents

Contents	ii
List of Figures	v
List of Tables	viii
1 Introduction	1
1.1 Motivation	1
1.2 Examples and Applications	3
1.3 The Tool: Cumulative Prospect Theory (CPT)	7
1.4 Related Work and Overview	13
Appendix	18
1.A Notational Conventions	18
2 Optimal Resource Allocation over Networks with CPT Players	20
2.1 Introduction	20
2.2 Lottery-Based Resource Allocation Model	21
2.3 Discretization Trick	23
2.4 Pricing and Market-Based Mechanism	25
2.5 A Quick Illustrative Example	29
2.6 Optimum Permutation Profile and Duality Gap	30
2.7 Average System Problem and Optimal Lottery Structure	33
2.8 Summary	36
Appendix	36
2.A Proof of Theorem 2.4.1	37
2.B Proof of Lemma 2.6.1	38
2.C Proof of Proposition 2.6.2	39
2.D Example to Show Duality Gap	40
2.E Proof of Theorem 2.6.3	43
2.F Proof of Lemma 2.7.1	44
2.G Proof of Lemma 2.7.2	45
2.H Proof of Proposition 2.7.3	47

3	Notions of Equilibrium: CPT Nash Equilibrium and CPT Correlated Equilibrium	49
3.1	Introduction	49
3.2	Definitions	50
3.3	Main Result: An Interesting Geometric Property	52
3.4	Two by Two Games	57
3.5	Connectedness	64
3.6	Example of a Game with Disconnected CPT Correlated Equilibrium	67
3.7	Summary	69
4	Black-Box Equilibrium: Reconsidering CPT Nash Equilibrium	71
4.1	Introduction	71
4.2	CPT and Betweenness	73
4.3	Equilibrium in Black-Box Strategies	77
4.4	Summary	85
	Appendix	87
4.A	Proof of Theorem 4.2.3	87
4.B	An Interesting Functional Equation	88
4.C	Proof of Lemma 4.3.6	90
4.D	Proof of Proposition 4.3.11	90
4.E	Proof of Theorem 4.3.12	90
4.F	Joint Continuity of the Concave Hull of a Jointly Continuous Function	92
5	Mediated Correlated Equilibrium: Reconsidering CPT Correlated Equilibrium	94
5.1	Introduction	94
5.2	Calibrated Learning in Games	96
5.3	Mediated Games and Equilibrium	100
5.4	Calibrated Learning to Mediated CPT Correlated Equilibrium	104
5.5	No-Regret Learning and CPT Correlated Equilibrium	106
5.6	Summary	111
	Appendix	111
5.A	Notions of equilibrium	111
5.B	Beyond Fixed Reference Points	113
5.C	Generalized Signal Spaces	116
5.D	Proof of Lemma 5.3.3	118
5.E	Proof of Theorem 5.4.1	121
5.F	Proof of Corollary 5.4.3	122
5.G	Proof of Proposition 5.4.4	122
5.H	Proof of Proposition 5.4.5	123
5.I	Proof of proposition 5.5.1	127
5.J	Proof of Lemma 5.5.4	128

5.K Proof of Proposition 5.5.3	128
6 Mediated Mechanism Design for CPT Players	136
6.1 Introduction	136
6.2 Mechanism Design Framework	141
6.3 The Revelation Principle	148
6.4 Mediated Mechanisms and the Revelation Principle	157
6.5 Summary	169
Appendix	172
6.A Proof of the Revelation Principle	172
6.B Proof of Equation (6.A.6)	180
6.C Outcome Sets can be Identified with the Allocation Set under EUT	183
7 Concluding Remarks and Directions for Future Work	187
7.1 Introduction	187
7.2 Role of Communication, Data Analytics and Artificial Intelligence in Resource Allocation	190
7.3 Fairness and Ethical Considerations	194
7.4 Conclusion	196
Bibliography	197
Index	206

List of Figures

- | | | |
|-----|--|----|
| 1.1 | Allais Paradox: Lotteries involved in the thought experiments proposed by Allais are shown. Each lottery is comprised of the winning amounts and the corresponding chance of winning these amounts. | 8 |
| 1.2 | Example of a typical value function. The plot shows value function v^0 (i.e. reference point $r = 0$) given in equation (1.3.5) with $\alpha_1 = \alpha_2 = 0.5$, and $\lambda = 2.5$. With reference point $r = 0$, positive outcomes ($z > 0$) are gains and negative outcomes ($z < 0$) are losses. The value at reference point $r = 0$ is 0. Notice that the value function is concave in the positive domain and convex in the negative domain. Also, notice that value function is much more steeper in the negative domain than in the positive domain giving rise to a kink at the origin. | 10 |
| 1.3 | Example of a typical probability weighting function. The solid curve shows the typical shape of a probability weighting function, be it for gains or for losses. The dotted line shows the identity function for reference. This marks the deviation from EUT. Indeed, if the probability weighting function is given by the identity function for both gains and losses, then the player has EUT preferences with the utility function given by the value function v^r at its reference point. The plot shows the probability weighting function given by equation (1.3.6) with $\gamma = 0.65$. Notice that the probability weighting function is typically concave initially and convex later with an inflection point around $1/3$. It over-weights smaller probabilities and under-weights larger probabilities. The probabilistic sensitivity (derivative of the probability weighting function) is high near the end probabilities, namely, 0 and 1, and low in the middle. | 11 |
| 1.4 | Agent-network decomposition: The central optimizer broadcasts a menu of market prices and the autonomous agents respond with the optimal budget signals based on the preferences learned from the users. These budget signals are then collected by the central optimizer to update the broadcasted prices and allocate resources. See Chapter 2 for more details. | 14 |
| 1.5 | Flow chart showing the dependence between different chapters and sections. . . | 17 |
| 2.1 | Probability weighting function for the example in Section 2.5. The plot shows the probability weighting function given by equation (1.3.6) (also shown in Section 2.5) with $\gamma = 0.61$. The dotted line shows the identity function for reference. | 29 |

3.1	Payoff matrix (left) and joint probability matrix (right) of a 2×2 game	57
3.2	Canonical 2×2 games	61
3.3	Vertices of the convex polytope C_{EUT} for $\gamma_I(\alpha, \beta)$	62
3.4	The set C_{CPT} for games of type (I) with weakly dominated strategies	63
3.5	The set C_{CPT} for games of type (II) with weakly dominated strategies	64
3.1	Standard 2-simplex of probability vectors $p = (p_R, p_Y, p_G)$. The shaded region represents the set $C(\Gamma, 1, \text{TOP})$ and is disconnected.	69
3.2	Un-normalized distributions $\bar{\mu}$ and $\tilde{\mu}$	70
4.1	The solid curve shows the probability weighting function for Alice from Example 4.2.1 and Example 4.3.7, and the dashed curve shows the probability weighting function for Charlie from Example 4.2.2.	75
4.1	Payoff matrix for Alice in Example 4.3.7. Rows and columns correspond to Alice's and Bob's actions respectively. The amount in each cell corresponds to Alice's payoff.	79
4.2	Payoff matrices for the 2×2 game in Example 4.3.9 (left matrix for player 1 and right matrix for player 2). The rows and the columns correspond to the actions of player 1 and player 2, respectively, and the entries in the cell represent the corresponding payoffs.	80
4.3	The CPT value of player 1 in Example 4.3.9. Here, p and q denote the black-box strategies for player 1 and 2, respectively. Note the rise and sharp drop in the two curves near $p = 1$. For the curve for $q = 0.3$, the global maximum is attained at $p = 0$, whereas, for the curve for $q = 0.35$, the global maximum is attained close to $p = 1$, specifically for some $p \in [0.9, 1]$	81
4.4	The figure (not to scale) shows the best response black-box strategy sets of the two players for the game in Example 4.3.9. The red (dashed) line shows the best response black-box strategy set of player 2 in response to the black-box strategy $(1 - p, p)$ of player 1. The green (solid) line shows the best response black-box strategy set of player 1 in response to the black-box strategy $(1 - q, q)$ of player 2. Note that there is no intersection of these lines.	82
4.5	Venn diagram depicting the different notions of equilibrium as subsets of the set $S = \prod_i \Delta(A_i)$. The sets marked pNE, mNE, BBNE, and mBBNE represent the sets of pure Nash equilibria, mixed action Nash equilibria, black-box strategy Nash equilibria, and mixed black-box strategy Nash equilibria, respectively. Examples are given in the body of the text of CPT games lying in each of the indicated regions (a) through (g).	84
4.6	Payoff matrices for the 2×2 games in Example 4.3.15. The rows and the columns correspond to the actions of player 1 and player 2, respectively. In each cell, the left and right entries correspond to player 1 and player 2, respectively. The labels indicate the corresponding regions in Figure 4.5. The game matrix for the example corresponding to region (g) is the same as that for the one corresponding to region (a).	86

4.7	The CPT value function for player 1 in Example 4.3.15(b), when $q = 0.5$ is the mixture of actions of player 2.	86
4.8	The CPT value function for player 1 in Example 4.3.15(e), when $q = 0.5$ is the mixture of actions of player 2.	86
6.1	The solid curve shows the probability weighting function w_1^+ for player 1 from Example 6.2.1 and Example 6.4.2. The dotted curve shows the probability weighting function w_2^+ for player 2 in Example 6.2.1 and Example 6.4.2, which is the linear function corresponding to EUT preferences.	147
6.1	Probability weighting functions for the players in Example 6.3.1.	150
6.2	Plot of expression $E(x)$ in Claim 6.3.2.	152
6.3	Plot of expression $E_1(z')$ in Claim 6.3.3.	153
6.4	Plot of expression $E_2(x')$ in Claim 6.3.3.	153
6.1	Plot of expression $E_3(x)$ in Example 6.4.3.	166

List of Tables

2.D.1	Table showing the numerical evaluations corresponding to the example in Section 2.D. The numbers in the cells denote the optimum value of the objective function (VIII) in the corresponding cases.	42
3.1	Payoff matrix for the game in Section 3.6	67
4.1	Different types of Nash equilibrium when players have CPT preferences.	87
5.1	Payoff matrix for the game Γ^* in Example 5.2.1. The rows and columns correspond to player 1 and 2's actions respectively. The first entry in each cell corresponds to player 1's payoff and second to player 2's payoff.	98
5.2	Empirical distribution μ^o for the action play in Example 5.2.1.	99
5.3	Empirical distribution μ^e for the action play in Example 5.2.1.	99
5.4	Empirical distribution μ^* for the action play in Example 5.2.1.	99
5.K.1	Empirical distribution $\hat{\mu}$ in Example 5.5.2.	131
6.1	Settings in which the revelation principle holds.	161

Acknowledgments

Firstly, I would like to thank my advisor Venkat Anantharam. I am grateful to him for pointing me towards cumulative prospect theory (CPT), which has taken center stage in my research and this thesis. His guidance has been pivotal in shaping my research. One of his amazing characteristics is that he is ultra-careful in everything he does, and that has greatly influenced my thinking, research, writing, and presentation. I am indebted to him for his detailed feedback on all the documents and for all our discussions, full of originality, precision, witticism, and enlightening thoughts. I would also like to thank Jean Walrand and David Aldous for being a part of my dissertation committee, and Anant Sahai for being a part of my qualifying exam committee. I am grateful to them for their helpful comments.

I thank my collaborators, Kannan Ramchandran, Thomas Courtade, Vipul Gupta, Vidya Muthukumar, and Anant Sahai. Conversations and correspondences with them have provided me with the much-needed boost that makes research an exciting social engagement. I am grateful to the wonderful conferences as Allerton, GameNets, and Stony Brook Game Theory Center for providing me many more such opportunities with world-class researchers. I am thankful to Inderjit Dhillon for allowing me to intern at Amazon. I would also like to thank my fellow researchers at Amazon, Daniel Hill, Kedarnath Kolluri, Arya Mazumdar, Sujay Sanghvi, and Rajat Sen, for the illuminating discussions and for helping me understand the intricacies involved in implementing algorithms at such a large scale. I had the opportunity to teach as a graduate student instructor at UC Berkeley for EE226 with Venkat Anantharam and CS70 with Alistair Sinclair and Yun Song. Working with them has been a great learning experience and has helped me improve my teaching and managerial skills. I am grateful to them for giving me this opportunity. I am deeply indebted to Vivek Borkar for advising me during my undergraduate studies and prodding me to join UC Berkeley. He has always been a great guide and a source of inspiration. I would also like to thank my IMO and IIT JEE coach, M. Prakash, who has advised me at every prime juncture of my career. One of the books that has helped me a lot during my PhD research is “Prospect Theory: For Risk And Ambiguity” by Peter Wakker, and I am grateful to him for writing this valuable piece of work.

Secondly, I would like to thank my friends and colleagues at Berkeley who helped make this journey memorable. BLISS lab has been an amazing place to be, full of intelligent people and exciting activities. I would like to thank Raaz Dwivedi, Sri Krishna Vadlamani, Sang Min Han, Kush Bhatia, Neha Gupta, Payam Delgosha, Vidya Muthukumar, Ashwin Pananjady, Orhan Ocal, Vijay Kamble, Varun Jog, Kangwook Lee, Nihar Shah, Rashmi Vinayak, Vasuki Narasimha Swamy, Fanny Yang, Yuting Wei, Dong Yin, Avishek Ghosh, Vipul Gupta, Aviva Bulow, Efe Aras, Kuan-Yun Lee, Vignesh Subramanian, Nived Rajaraman, Wenlong Mou, Banghua Zhou, Chinmay Maheshwari, Swanand Kadhe, Abishek Sankararaman, Gireeja Ranade, Ramtin Pedarsani, Kabir Chandrashekher, Leah Dickstein, Tavor Baharav, Laura Brink, Aditya Ramdas, Reinhard Heckel, Ilan Shomorony, Ludwig Schmidt, Koulik Khamaru, Milind Hegde, Anamika Chowdhury, Jimmy Narang, Nick Bhattacharya, Satyaki Mukherjee. In particular, thanks to Raaz and Sri Krishna for taking

this journey together. Thanks to Kush and Neha for being a part of my pandemic bubble and helping me keep my sanity during these crazy times. I owe thanks to Shirley Salanio, Kim Kail, and Nicole Song for being extremely reliable and resourceful in resolving all my administrative issues.

Finally, I would like to thank my parents – my mother Archana Phade, and my father, Rajesh Phade. Without their love and support none of this would have been possible, and this thesis is dedicated to them.

Chapter 1

Introduction

1.1 Motivation

We will mainly be concerned with the study of social systems comprised of several individuals, typically humans, henceforth called *players*, interacting directly or indirectly in a bounded situation (or an environment). Systems influenced by technological innovations over the past several decades will be of particular interest to us. For example, these include transportation and communication networks, the Internet, computation networks and data-centers, energy and utility networks, financial networks, labor markets, social networks, and digital markets.

The complex nature of these systems requires consideration of several crucial aspects which gave rise to the interdisciplinary fields of *cybernetics* and *systems science*. These combine knowledge from various fields such as control theory, information theory, dynamical systems, operations research, computer science, systems engineering, economics, statistics, and psychology. The engineering approach towards solving these problems primarily focuses on the physical aspects such as feasibility, practicality, maintainability, stability, and scalability. An equally important dimension is that of catering to individual preferences and needs. Ultimately these systems are there for the users. Thus enters marketing research and business management. These fields study the market economy and business processes to identify, anticipate and satisfy customers' needs and wants. A holistic approach that combines these two approaches will go a long way.

Technological advancements in domains such as the Internet, Computing, Communication, and Artificial Intelligence (AI) have lead to rapidly evolving network services such as cloud computing, smart information systems, multimedia platforms, software companies, online marketplaces, and smart grids, that have global scopes. Consequently, network economics research evolved along two major lines:

1. Optimal routing and control: This involved the study of flow dynamics and congestion based on the underlying network structures and routing decisions. Typical problems studied include the shortest path problem, the maximum flow problem, the minimum cost flow problem, etc. (See books by Anna Nagurney [94, 95, 97, 98, 96].)

2. Network formation and growth: Here, the focus is on the understanding of the formation of network links, the flow of information in social networks or diseases in epidemiological studies, connectivity and segregation in different networks, etc. Models from random graph theory and statistics are helpful in this approach. (See books by Mathew Jackson [62, 63] and Sanjeev Goyal [54].)

Besides understanding the working of networks, a fundamental goal of network economics is to *assist decision-making* for both the system designer and the players in the system. For example, Braess' paradox warns a network planner of the following counter-intuitive effect: adding additional links to a network can reduce the overall system utility (such as the total delays for all the drivers in a transportation network) at Nash equilibrium when each player is making an optimal self-interested decision. Observations like these and results from network economics have greatly helped policy-making and system design. (Shapiro and Varian [122] describe strategies to guide business decisions and policies in network economies such as differential pricing, utilizing network positive externalities and lock-in effects, patents and rights management, and others.)

Game theory and economics offer valuable guiding principles in the design of these systems. The economic models for studying these problems typically assume that the participating agents are rational and possess immense computational power (which is reasonable when the participating agents are firms or nations). However, for e-commerce platforms like social media and online marketplaces, where the participating agents are *single individuals* who perform *several repeated short-lived interactions* with the platform, it is unusual that these agents would adhere to the above behavioral assumptions. We cannot expect the human mind to make informed and well-thought decisions in such complex interconnected systems, let alone the stress it generates. Our goal here is to use sophisticated models from behavioral psychology and decision theory to model human interaction and design robust and scalable systems that would assist the users in making decisions that are in their own interests and also for those around them.

The digital revolution has given rise to software companies having massive control over several crucial networks with the power to micromanage them. The algorithms deployed by these companies can influence social, economic, and political networks like never before. Along with all the evident benefits of these software systems in automating tasks and facilitating large-scale network operations, we must pay closer attention to how these systems interact with their users. The growing human-computer interaction requires careful consideration of human behavior and their emotional responses. Our knowledge regarding the guiding principles for governing these interactions is quite limited, and a methodological approach towards incorporating psychological aspects into system design is barely off the ground. There is an ongoing debate relating to the benefits of these big technology companies, the extreme power these companies hold, and whether they are using it wisely or not. Although it will not be the focus of this thesis, I hope that the behavioral foundations developed in this work would help answer some of these questions (see Section 7.3), and consequently, help build systems that are better aware of human behavior and needs.

Perhaps the most apt historical model for algorithmic regulation is not trust-busting, but environmental protection. To improve the ecology around a river, it isn't enough to simply regulate companies pollution. Nor will it help to just break up the polluting companies. You need to think about how the river is used by citizens—what sort of residential buildings are constructed along the banks, what is transported up and down the river—and the fish that swim in the water. Fishermen, yachtsmen, ecologists, property developers, and area residents all need a say. Apply that metaphor to the online world: Politicians, citizen-scientists, activists, and ordinary people will all have to work together to co-govern a technology whose impact is dependent on everyone's behavior, and that will be as integral to our lives and our economies as rivers once were to the emergence of early civilizations.

Anne Applebaum

“The Internet doesn't have to be awful.” *The Atlantic*. April 2021.

1.2 Examples and Applications

Transportation Networks

Let's say you want to reach the airport to catch a flight. You open a navigation app, such as Google Maps or Apple Maps, and check for possible routes and the estimated times of arrival. Your topmost concern is to arrive at your destination in time. Plus, you'd like to have a good estimate of your arrival time. Compare it with someone who might be using the same app but is looking for a scenic route and not so worried about his arrival time. At any given time, hundreds of thousands of users are using such apps to find what suits them the best. All these different people have varied requirements based on their purposes and preferences while sharing the same infrastructure and resources. The app recommendations affect their choices, and their choices have externalities that affect the conditions for others. One could imagine the app providing signals and economic incentives to alter traffic patterns.

A familiar example in this spirit is clearing the way for emergency vehicles. Something that we have been doing for several years. Another example is charging a variable rate adapted to the traffic conditions for the use of the express lanes. Given the prevalent use of navigation apps and other communicating devices today, we have more options to influence traffic routing. At the same time, we can collect and process a lot more data. Our goal is to explore ideas along these lines. An important thing to notice here is that the players in this system are human agents and they are bound to display behavioral features that do not fall under the traditional notions of rationality. For example, drivers might prefer routes that they are familiar with, even if the alternative route is faster. (This is reminiscent of the well-documented *endowment effect*, which says that people are more likely to hold onto an object they own rather than trade it for an equally or higher valued alternative they do not

own. The fear of the unknown and uncertainty also plays a role here.) We must incorporate these behavioral features into system modeling. Furthermore, this applies to all forms of transportation services such as public transport, railways, airways, waterways, shipping of goods, etc.

Communication Networks

Using navigation apps to help route traffic is just an instance of taking advantage of the advanced communication technologies for improving resource allocation. Indeed, communicating the availability of resources, individual preferences, and incentives for resource management, and controlling system parameters require real-time information transfer and signaling. No wonder the Internet was the first to witness real-time algorithm-based traffic management. Transmission Control Protocol (TCP) and bandwidth allocation algorithms have helped avoid the congestion issues that had plagued the Internet before TCP. The theoretical foundations for this work were laid by Kelly in the late 1990s [72, 73]. In Chapter 2, we extend these ideas to incorporate behavioral features and psychological traits displayed by the users.

Today, traffic shaping is a major area that deals with congestion control [84, 116]. The users are allocated bandwidth based on the choice of the monthly plans selected by them and the ambient network traffic conditions. One of our goals is to extend these ideas to real-time traffic management. For example, imagine you have a virtual presentation coming up. It would be nice to indicate this to the service provider, such as Xfinity or AT&T, and request a boost for this period. It might result in additional charges, but it would provide you the added benefit of choosing a more economical base plan. Certainly, re-engineering the Internet along these lines would increase user-system interactions and it would need algorithms that are more aware of human behavior and responses.

Cloud Computing Networks

Just as communication networks allocate bandwidth to the users, cloud computing networks, such as Amazon Web Services, Microsoft Azure, or Google Cloud, provide on-demand computer system resources such as data storage and computing power. Cloud service providers can schedule most of the customer jobs instantly today as the resources exceed the demand. However, with a growing trend of customers opting for computing resources as a service instead of maintaining such systems on their own, this surplus luxury is not sustainable. Resources are also naturally constrained in settings such as fog computing and peer-to-peer computing networks. Besides, concerns over the energy consumption by data centers is another factor that limits the expansion of computing resources.

The demand for resources can vary significantly over time, different jobs have different resource requirements, and customers have varying preferences towards their job delays and the quality of service. The prices must conform to these changing demands in real-time. Although the typical customers in this setting are firms and organizations, the end-users

of their services and products are often individual humans. The value and revenue generation for these organizations is closely related to the levels of consumer satisfaction. As a result, behavioral considerations naturally creep into the utilities and preferences of these organizations.

Energy Networks

Smart grids are another excellent example of the application of digital processing and communications to systems where user interactions play a major role. The goal here is to improve the economic efficiency of electricity networks and maintain high levels of quality of supply by integrating the behavior and actions of all the users connected to the network - generators, consumers, and those that do both. It would provide communication protocols to the suppliers and the consumers, allowing them to be more flexible and sophisticated in their operational strategies. For example, the suppliers could indicate their energy prices, and the consumers could indicate their willingness to pay in real-time. The users can configure smart devices to generate additional energy or initiate energy-saving modes under specific settings such as during high-cost peak usage periods. Similar to the pricing based on job delays in the cloud computing setting, we can imagine customers having different preferences towards their energy requirements based on deadlines, for example, such requirements would naturally occur in charging of electric vehicles. Today, PG&E, a utility company that provides natural gas and electric service, offers different pricing schemes such as time-of-use rate plans and tiered usage rate plans. Along similar lines, we are interested in much more flexible and sophisticated pricing schemes based on dynamic market conditions and human behavior analysis. This would also benefit in incentivizing people to shift to clean electricity options and adopt solar panels at home.

Social Networks

Several activities such as advertising, campaigning, or running welfare programs depend on the underlying social networks. Humans are the primary agents in any social network. Their interactions and behavior form an integral part in the study of social networks. Models that incorporate psychological aspects are needed to better allocate resources in these activities. It would help answer questions like: How can we maximize the impact of a campaign with a limited budget? How to best incentivize the agents in a network to perform actions that are in the best interests of the entire society?

From a commercial point of view, it would greatly benefit the online ad exchange companies such as Google Ads or Facebook Ads. These are digital marketplaces that enable advertisers to buy and sell advertising spaces. Here, user attention is the limited resource and the different advertisers are competing for this limited resource. The tools developed in this thesis will help regulate these markets more efficiently by incorporating human behavioral features.

Matching Markets

Just like the ad exchange marketplace, several other matching markets fall in this domain. These include labor markets that match employers and workers such as Upwork and Freelancer, ride hailing applications that match drivers and riders such as Uber and Lyft, delivery services that match restaurants and diners such as Doordash and UberEats, or online marketplaces that match sellers and buyers such as Amazon and eBay. Notice that most of the participating agents in these settings are individual humans susceptible to showing behavior that is influenced by biases and heuristics.

Finance and Insurance

Finance and insurance is another interesting setting where behavioral factors play a huge role. There is a decent amount of work studying how individuals make decisions about their investment strategies and insurance policies, but there is only a limited amount of work that considers behavioral features in a financial network setting where the individuals interact with each other and their decisions affect the other individuals in the network. In this work, we establish results that would facilitate this research.

Observe that, in all the above examples, the following factors are common:

1. The resources are limited.
2. Players have varying requirements and preferences.
3. The preferences of the players are private information.
4. Players have limited information about the system operations and constraints.
5. Players show behavioral features.

The goal is to design a communication protocol or a market system to facilitate the exchange of information for strategic players who might display behavioral features, and consequently allocate resources to satisfy certain requirements. In contrast to prior works, we will pay special attention to the last factor, namely, the behavioral features of the players. We aim to bring these aspects to the same level of mathematical sophistication as other aspects in system sciences. Such an approach is crucial to building systems that are scalable across different users and robust to the intricacies of human behavior.

With me, everything turns into
mathematics.

Rene Descartes

1.3 The Tool: Cumulative Prospect Theory (CPT)

Central to our approach is a mathematical model to capture human behavior and preferences. As is common in decision theory, we will consider the problem of decision-making by rational agents under uncertainty. In many of the examples discussed above, the agents need to make decisions without having complete information about the system and the behavior of other players in the system. For instance, a person who is traveling needs to decide which route to take without the exact information about traffic conditions, or a company launching a new product needs to decide how to maximize its advertising impact without complete knowledge of its customers as well as its competitors. Decision-making under uncertainty provides a minimal framework that is general enough to capture the commonly encountered interactions, preferences, choices and actions of agents in a network.

Rationality is generally formulated as expected utility maximization. The justification for this comes from the von Neumann and Morgenstern expected utility maximization theorem [130]. Although this assumption has a nice normative appeal to it and can be used to a large extent as a prescriptive theory, it has been evident through several examples [3, 48, 67] that the model is not that good an approximation for descriptive purposes. On the other hand, cumulative prospect theory (CPT) accommodates many empirically observed behavioral features [127]. Proposed by Kahneman and Tversky, it is one of the leading theories for decision making under uncertainty. It has a nice mathematical formulation and is a generalization of expected utility theory (EUT).

A *lottery* (or *prospect*) is comprised of one or more outcomes with their corresponding probabilities.¹ We will denote a lottery by

$$L := \{(p_1, z_1), (p_2, z_2), \dots, (p_t, z_t)\}, \quad (1.3.1)$$

where $z_j \in \mathbb{R}$, $1 \leq j \leq t$, denotes an *outcome* and p_j , $1 \leq j \leq t$, is the probability with which outcome z_j occurs. We assume that the lottery is *exhaustive*, i.e. $\sum_{j=1}^t p_j = 1$. (Note that we are allowed to have $p_j = 0$ for some values of j and we can have $z_k = z_l$ even when $k \neq l$.)

Expected utility theory (EUT) posits that each individual is associated with a utility function $u : \mathbb{R} \rightarrow \mathbb{R}$. The utility function is typically assumed to be concave (to capture the risk-averseness of the individual). The expected utility corresponding to lottery L is given by

$$U(L) = \sum_{j=1}^t p_j u(z_j). \quad (1.3.2)$$

A person is said to have EUT preferences if, given a choice between lottery L_1 and lottery L_2 , she chooses the one with the higher expected utility.

In the latter half of the 20th century, several people began documenting the limitations of EUT to model human behavior. Allais paradox (1952) is a particularly interesting thought experiment that marks the beginning of this work [3]. The experiment goes as follows: Consider the two lotteries shown in Experiment A of Figure 1.1. If you choose Lottery 1A, then you win \$1 Million for sure, i.e. with a 100% chance. If you choose Lottery 2A, then

Experiment A			
Lottery 1A		Lottery 2A	
Winning	Chance	Winning	Chance
\$1 Million	100%	\$1 Million	89%
		Nothing	1%
		\$5 Million	10%

Experiment B			
Lottery 1B		Lottery 2B	
Winning	Chance	Winning	Chance
Nothing	89%	Nothing	90%
\$1 Million	11%		
		\$5 Million	10%

Figure 1.1: Allais Paradox: Lotteries involved in the thought experiments proposed by Allais are shown. Each lottery is comprised of the winning amounts and the corresponding chance of winning these amounts.

you win \$1 Million with a 89% chance, you do not win anything with a 1% chance, and you win \$5 Million with a 10% chance. You can choose only one of the two lotteries and it is a one time offer. Which one do you select? Now instead, consider Experiment B. The lotteries are shown in Figure 1.1. Again you can choose only one of the two lotteries and it is a one time offer. Which one now?

It is quite common for people to choose Lottery 1A in Experiment A and Lottery 2B in Experiment B. The paradox arises from the following observation: Lottery 1B is obtained from Lottery 1A by transforming 89% chance of winning \$1 Million to winning nothing. The remaining 11% chance of winning \$1 Million is left as is. When a similar transformation is applied to Lottery 2A, we get Lottery 2B. Note that the 89% chance of winning \$1 Million in Lottery 2A changed to nothing combined with the 1% chance of winning nothing gives the 90% chance of winning nothing in Lottery 2B. The 10% chance of winning \$5 Million remains unchanged.

The above observation underlies the fact that the choice of Lottery 1A and Lottery 2B is inconsistent with EUT. To see this, let u be the utility function of the individual. Choice of Lottery 1A over Lottery 2A implies

$$u(\$1M) > 0.89u(\$1M) + 0.01u(\$0) + 0.10u(\$5M). \quad (1.3.3)$$

And choice of Lottery 2B over Lottery 1B implies

$$0.89u(\$0) + 0.11u(\$1M) > 0.90u(\$0) + 0.10u(\$1M). \quad (1.3.4)$$

No utility function u can satisfy the above inequalities simultaneously. Later we will see how CPT explains this phenomenon.

We now give a quick review of *cumulative prospect theory* (CPT) (for more details see [132]). Each person is associated with a *reference point* $r \in \mathbb{R}$, a corresponding *value function* $v^r : \mathbb{R} \rightarrow \mathbb{R}$, and two *probability weighting functions* $w^\pm : [0, 1] \rightarrow [0, 1]$, w^+ for gains and w^- for losses. We say that (r, v^r, w^\pm) are the *CPT features* of that person.

The function $v^r(x)$ satisfies: (i) it is continuous in x ; (ii) $v^r(r) = 0$; (iii) it is strictly increasing in x . The value function is generally assumed to be convex in the losses domain ($x < r$) and concave in the gains domain ($x \geq r$), and to be steeper for losses for gains in the sense that $v^r(r) - v^r(r - z) \geq v^r(r + z) - v^r(r)$ for all $z \geq 0$.² The reference point is meant to capture psychological factors such as the player's expectations, her status quo, or her goal. By letting the value function be steeper for losses we are able to capture the individual's loss aversion. The concavity in gains and convexity in losses captures the effect of diminishing sensitivity of the individual. Contrast this with the typical assumption that the utility function is concave throughout in EUT.

An example of a typical value function is

$$v^r(z) = \begin{cases} (z - r)^{\alpha_1} & \text{for } z \geq r, \\ -\lambda(r - z)^{\alpha_2} & \text{for } z < 0, \end{cases} \quad (1.3.5)$$

with $\alpha_1, \alpha_2 \in (0, 1]$, and $\lambda \geq 1$. Here, α_1 and α_2 capture the diminishing sensitivity to returns for gains and losses, respectively, and λ captures the loss aversion. In Figure 1.2, we plot the above value function with $\alpha_1 = \alpha_2 = 0.5$, $\lambda = 2.5$.

The probability weighting function along with the ordering of the outcomes in a lottery, dictates the *probabilistic sensitivity* of a player, a property that plays an important role in lotteries and gambling. As Boyce [3] points out, "It is the lure of getting the good without having to pay for it that gives allocation by lottery its appeal."

The probability weighting function typically over-weights small probabilities and under-weights large probabilities, and this captures the 'lure' effect. The probability weighting functions $w^\pm : [0, 1] \rightarrow [0, 1]$ satisfy: (i) they are continuous; (ii) they are strictly increasing; (iii) $w^\pm(0) = 0$ and $w^\pm(1) = 1$.

An example of a typical probability weighting function (for gains or losses) suggested by Prelec [113] is

$$w(p) = \exp\{-(-\ln p)^\gamma\}, \quad (1.3.6)$$

with $\gamma \in (0, 1]$. In Figure 1.3, we plot this function with $\gamma = 0.65$.

We now describe how to compute the CPT value of a lottery. This is the analog of the expected utility in EUT. Let $\alpha := (\alpha_1, \dots, \alpha_t)$ be a permutation of $(1, \dots, t)$ such that

$$z_{\alpha_1} \geq z_{\alpha_2} \geq \dots \geq z_{\alpha_t}. \quad (1.3.7)$$

Let $0 \leq j_r \leq t$ be such that $z_{\alpha_j} \geq r$ for $1 \leq j \leq j_r$ and $z_{\alpha_j} < r$ for $j_r < j \leq t$. (Here $j_r = 0$ when $z_{\alpha_j} < r$ for all $1 \leq j \leq t$.) The *CPT value* $V^r(L)$ of the prospect L is evaluated using the value function $v^r(\cdot)$ and the probability weighting functions $w^\pm(\cdot)$ as follows:

$$V^r(L) := \sum_{j=1}^{j_r} \nabla_j^+(p, \alpha) v^r(z_{\alpha_j}) + \sum_{j=j_r+1}^t \nabla_j^-(p, \alpha) v^r(z_{\alpha_j}), \quad (1.3.8)$$




Figure 1.2: Example of a typical value function. The plot shows value function v^0 (i.e. reference point $r = 0$) given in equation (1.3.5) with $\alpha_1 = \alpha_2 = 0.5$, and $\lambda = 2.5$. With reference point $r = 0$, positive outcomes ($z > 0$) are gains and negative outcomes ($z < 0$) are losses. The value at reference point $r = 0$ is 0. Notice that the value function is concave in the positive domain and convex in the negative domain. Also, notice that value function is much more steeper in the negative domain than in the positive domain giving rise to a kink at the origin.

where $\nabla_j^+(p, \alpha)$, $1 \leq j \leq j_r$, $\nabla_j^-(p, \alpha)$, $j_r < j \leq t$, are *decision weights* defined via:

$$\begin{aligned} \nabla_1^+(p, \alpha) &:= w^+(p_{\alpha_1}), \\ \nabla_j^+(p, \alpha) &:= w^+(p_{\alpha_1} + \dots + p_{\alpha_j}) - w^+(p_{\alpha_1} + \dots + p_{\alpha_{j-1}}) && \text{for } 1 < j \leq t, \\ \nabla_j^-(p, \alpha) &:= w^-(p_{\alpha_t} + \dots + p_{\alpha_j}) - w^-(p_{\alpha_t} + \dots + p_{\alpha_{j+1}}) && \text{for } 1 \leq j < t, \\ \nabla_t^-(p, \alpha) &:= w^-(p_{\alpha_t}). \end{aligned}$$

Although the expression on the right in equation (1.3.8) depends on the permutation α , one can check that the formula evaluates to the same value $V^r(L)$ as long as the permutation α




Figure 1.3: Example of a typical probability weighting function. The solid curve shows the typical shape of a probability weighting function, be it for gains or for losses. The dotted line shows the identity function for reference. This marks the deviation from EUT. Indeed, if the probability weighting function is given by the identity function for both gains and losses, then the player has EUT preferences with the utility function given by the value function v^r at its reference point. The plot shows the probability weighting function given by equation (1.3.6) with $\gamma = 0.65$. Notice that the probability weighting function is typically concave initially and convex later with an inflection point around $1/3$. It over-weights smaller probabilities and under-weights larger probabilities. The probabilistic sensitivity (derivative of the probability weighting function) is high near the end probabilities, namely, 0 and 1, and low in the middle.

satisfies (1.3.7). The CPT value in equation (1.3.8) can equivalently be written as:

$$\begin{aligned}
 V^r(L) &= \sum_{j=1}^{j_r-1} w^+ \left(\sum_{i=1}^j p_{\alpha_i} \right) [v^r(z_{\alpha_j}) - v^r(z_{\alpha_{j+1}})] \\
 &\quad + w^+ \left(\sum_{i=1}^{j_r} p_{\alpha_i} \right) v^r(z_{\alpha_{j_r}}) + w^- \left(\sum_{i=j_r+1}^t p_{\alpha_i} \right) v^r(z_{\alpha_{j_r+1}}) \\
 &\quad + \sum_{j=j_r+1}^{t-1} w^- \left(\sum_{i=j+1}^t p_{\alpha_i} \right) [v^r(z_{\alpha_{j+1}}) - v^r(z_{\alpha_j})]. \tag{1.3.9}
 \end{aligned}$$

A person is said to have CPT preferences if, given a choice between prospect L_1 and prospect L_2 , she chooses the one with higher CPT value.

Let us see how CPT explains the Allais paradox. Let the value function be given by

$$v^r(z) = \begin{cases} (z - r)^{0.5} & \text{for } z \geq r, \\ -10(r - z)^{0.5} & \text{for } z < r. \end{cases}$$

Interpret this as the value of winning \$z Million is $v^r(z)$. Let the probability weighting functions for both gains and losses be as shown in Figure 1.3. When faced with the two lotteries in Experiment A, suppose the reference point is \$1 Million, i.e. $r = 1$. The CPT value of Lottery 1A is 0 (recall that $v^r(r) = 0$). The CPT value of Lottery 2A is given by

$$w^-(0.01)v^1(0) + w^+(0.1)v^1(5) = 0.0673 \times (-10) + 0.1791 \times 2 = -0.3148.$$

Hence, Lottery 1A is preferred over Lottery 2A. When faced with the two lotteries in Experiment B, suppose the reference point is 0, i.e. $r = 0$. Then, the CPT value of lottery 1B is given by

$$w^+(0.11)v^0(1) = 0.1877,$$

and the CPT value of Lottery 2B is given by

$$w^+(0.1)v^0(5) = 0.1791 \times 2.2361 = 0.4005.$$

Thus, Lottery 2B is preferred over Lottery 1B. This resolves the Allais paradox. Notice how the different aspects in CPT play a role here: the reference point captures the expectations of the individual and the 1% chance of winning nothing in Lottery 2A is thus perceived as a loss. Combined with the over-weighting on the 1% chance by the probability weighting function w^- and the high loss aversion of $\lambda = 10$ makes Lottery 2A disfavored as compared to Lottery 1A. On the other hand, in Experiment B, we assumed the reference point to be 0. The probability weighting function w^+ assigns similar decision weights to winning \$1 Million and \$5 Million (namely, $w^+(0.11)$ and $w^+(0.1)$, respectively). Naturally, winning \$5 Million is favored over winning \$1 Million, and Lottery 2B is preferred over Lottery 1B.

In a way, CPT seems to introduce so much flexibility that one could fit almost any observation. As said by John von Neumann, “with four parameters I can fit an elephant, and with five I can make him wiggle his trunk.” In this regard, I would like to point out that each of the concepts introduced by CPT such as the reference point, value function, and probability weighting functions have an interpretation that fulfills certain behavioral requirements. Moreover, this flexibility allows CPT to encompass a plethora of behavioral aspects in a convenient mathematical form. Finally, and most importantly, any theoretical guarantees provided under such flexible settings are applicable in restricted settings of CPT and hence not affected by the potential overparametrization present in CPT.

CPT also satisfies some important properties such as:

- *Strict stochastic dominance* [30]: shifting positive probability mass from an outcome to a strictly preferred outcome leads to a strictly preferred prospect. For example, the prospect $L_1 = \{(0.6, 40); (0.4, 20)\}$ can be obtained from the prospect $L_2 =$

$\{(0.5, 40); (0.5, 20)\}$ by shifting a probability mass of 0.1 from outcome 20 to a strictly better outcome 40. The strict stochastic dominance condition says that $V^r(L_1) > V^r(L_2)$ (see equation (1.3.9)).

- *Strict monotonicity* [30]: any prospect becomes strictly better as soon as one of its outcomes is strictly improved. For example, if $L_1 = \{(0.6, 40); (0.4, -10)\}$ and $L_2 = \{(0.6, 40); (0.4, -20)\}$, then $V^r(L_1) > V^r(L_2)$ (see equation (1.3.8)).

One wonders whether it is necessary to have the cumulative form of probability weighting in the evaluation of the CPT value. In fact, a precursor to cumulative prospect theory was proposed by Kahneman and Tversky in 1979, called *prospect theory* (PT). However, in the subsequent years, several drawbacks of this theory were observed. For example, it does not satisfy the first order stochastic dominance property. Schmedler[119] and Quiggin[114] developed rank dependent utility theory (RDU) that involved the cumulative functional form. CPT combines RDU with reference dependence in a consistent manner. Wakker, in his book [132], argues how the cumulative functional form is indeed the correct way to extend EUT to account for probabilistic sensitivity.

To summarize, CPT has the following features that are lacking from EUT: (i) reference dependence, (ii) loss aversion, (iii) diminishing sensitivity to returns for both, gains and losses, (iv) probabilistic sensitivity, (v) rank dependence and cumulative probability weighting.

1.4 Related Work and Overview

Several works have confirmed the applicability of prospect theory and cumulative prospect theory to individual decision-making in laboratory settings [66]. There is also evidence that these theories offer good description of behavior for the participants in game shows with large prizes [65, 110] and professional investors in financial markets [1]. The application of these theories to economics is relatively scarce. Amongst those, application of prospect theory to the fields of finance and insurance is by far most popular [14, 11, 39, 9, 42] (see also, [10] and the references therein). Most of these studies are lacking along two major lines:

1. they consider prospect theory to model behavior instead of its improved version, namely, cumulative prospect theory, and
2. they restrict their attention to single-agent choice scenarios without much consideration for the effects arising from multi-agent interactions.

In this thesis, we will develop theory that contributes to both these directions. To quote Camerer [24],

There is no good scientific reason why it (prospect theory or rather cumulative prospect theory) should not replace expected utility in current research, and be given prominent space in economics textbooks.

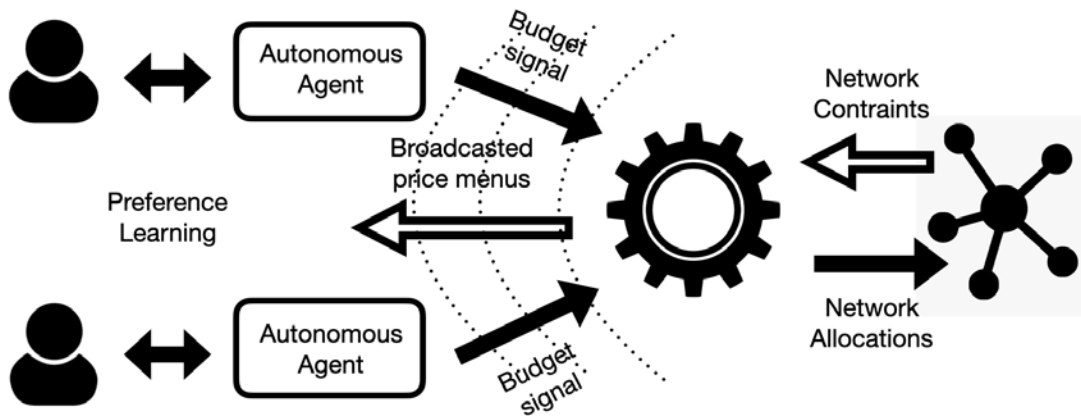


Figure 1.4: Agent-network decomposition: The central optimizer broadcasts a menu of market prices and the autonomous agents respond with the optimal budget signals based on the preferences learned from the users. These budget signals are then collected by the central optimizer to update the broadcasted prices and allocate resources. See Chapter 2 for more details.

As observed above in the examples, resource allocation over networks is a fundamental problem in systems. In Chapter 2, we consider this problem with CPT players. Modeling the agents using CPT brings forward the surprising benefits of lottery-based allocations in increasing social welfare compared to the corresponding EUT-based analysis. This conforms with the observation that lotteries help in influencing the behavior of people, which is backed by several experimental studies [111]. Our work³ based on CPT analysis not only explains the above phenomenon theoretically but also provides a practical algorithm to design the optimum lottery allocations. It achieves this without actually knowing the CPT preferences of the agents through an appropriate pricing scheme.

The practicality of this algorithm stems from the *agent-network decomposition*, outlined in Figure 1.4. It consists of decomposing the network resource allocation problem into a central optimizer and several user optimizers, one for each user. It builds on the work by Kelly [72] on optimal bandwidth allocation. Similar to the TCP/IP used in internet routing, a signaling scheme is set up between the user optimizers and the central optimizer. Here the central optimizer broadcasts a menu of prices, and the user optimizers respond with their optimal budget signals based on their individual preferences. We show that such an iterative scheme converges to the optimal lottery scheme and can track the trends in gradually changing market behavior. To accommodate CPT preferences, we invent a novel pricing scheme. It is actually reminiscent of the pricing strategies employed by several airline companies where they offer their customers an option to make additional payments to be enrolled in a raffle to win an upgrade on their tickets.

In this analysis, each agent is assumed to be a price-taker, i.e., we assume that the agents respond in a myopically optimal manner to the broadcasted prices, and none of the agents are dominant enough to single-handedly influence the market prices. However, when there are localized interactions amongst a small number of agents, game-theoretic models become important to study these interactions. The popular choice for studying these models is through different notions of equilibrium.

In Chapter 3, we take the notions of CPT Nash equilibrium and CPT correlated equilibrium defined by Keskin [74] as our starting point and establish several geometric properties of these equilibrium notions.⁴ We explore questions such as the convexity and the connectivity properties of the correlated equilibria set, and the relation between the Nash equilibria and the correlated equilibria. For example, under EUT, it was known that the set of all correlated equilibria is a convex polytope. Keskin showed that this property need not hold under CPT. We prove that it can, in fact, be disconnected. Nonetheless, certain properties like the Nash equilibria all lying on the boundary of the correlated equilibria set [101] continue to hold true, although they require new proof techniques.

These new theoretical phenomena bring out some of the important distinctions resulting from CPT modeling. One of the fundamental reasons for these distinctions is that CPT preferences do not satisfy something called the betweenness property. Betweenness implies that if an agent is indifferent between L_1 and L_2 , then she is indifferent between any mixtures of them too. Several empirical studies show systematic violations of betweenness [25, 2, 41, 123], and this makes the use of CPT more attractive than EUT for modeling human preferences. Further evidence comes from [25], where the authors fit data from nine studies using three non-EUT models, one of them being CPT, to find that, compared to the EUT model, the non-EUT models perform better.

Building upon the idea that the players might prefer to actively randomize their actions, we consider mixed strategies in non-cooperative games from a new perspective. We refer to such actively mixed strategies as black-box strategies.

Traditionally, mixed actions have been considered from two viewpoints, especially in the context of mixed-action Nash equilibrium. According to the first viewpoint, these are conscious randomizations by the players—each player only knows her mixed-action and not its pure realization. The notion of black-box strategies captures this interpretation of mixed-actions. According to the other viewpoint, players do not randomize, and each player chooses some definite action. But the other players need not know which one and the mixture represents their uncertainty, i.e., their conjecture about her choice.

Under CPT, these two interpretations get nicely untangled, and we get four different concepts of Nash equilibria depending on whether we allow randomization in each of the interpretations. In Chapter 4, we develop these four notions and study their properties such as existence and relation to each other.⁵

We then consider the setting of learning in repeated games in Chapter 5.⁶ The literature on learning in games provides an alternative explanation to the equilibrium notions as a long-run outcome in repeated games with mild rationality assumptions on the players. They are especially important from a behavioral perspective where players have limited computational

powers. We consider the celebrated theorem of Vohra and Foster [50] on the convergence of the empirical average of the action play to the set of correlated equilibria when players make calibrated forecasts and respond with myopically optimal actions. One soon realizes that the notion of CPT correlated equilibrium, as defined by Keskin, is not enough to capture this result. But instead, we need an appropriate convexification of this set that we call the mediated CPT correlated equilibrium.

In a correlated equilibrium, the mediator is assumed to recommend an action to each player to play. We introduce the notion of mediated games in which the mediator is allowed to send signals from more general sets. This is a specific type of game with communication as introduced by Myerson. The mediated CPT correlated equilibria are then the Bayes-Nash equilibria of this mediated game. Since the mediated CPT correlated equilibria are more general than the CPT correlated equilibria we get that the revelation principle in the context of correlated equilibria does not hold under CPT preferences.

Calibrated learning is one form of learning in games. More generally, the result on the convergence to correlated equilibria is closely related to the notion of no-regret learning in games. We prove that the set of CPT correlated equilibria is not approachable in the Blackwell approachability sense. These results strongly suggest that the notion of mediated CPT correlated equilibrium is the appropriate notion to consider in this context.

A major revelation in the previous study is that: The revelation principle fails under CPT! A natural question is what happens in mechanism design where the revelation principle has played a fundamental role. As suspected we observe that the revelation principle fails in mechanism design when agents have CPT preferences. In Chapter 6, we develop an appropriate framework that we call mediated mechanism design that allows us to recover the revelation principle under certain settings.⁷

The premise of a typical mechanism design scenario comprises a bunch of agents each having a private type consisting of their private information and preferences over the outcomes. There is a system operator (or a principal) who controls the implementations in the system but cannot directly observe the private types of players. To achieve optimal implementations conditioned on the types of the players, the system operator designs a communication protocol where the players can interact strategically in the resulting game. This allows the system operator to elicit information about the private types of the players. As an example, think of auctions. The second-price sealed-bid auctions incentivize the players to reveal their private values truthfully, and the item is allocated to the player with the highest value at the second-highest bid.

For a modern application, consider ride-hailing services such as Uber or Lyft. These apps present the customers with several options such as premium rides, shared rides, economy rides, etc. The purpose of these options is to elicit the preferences of the customers and provide optimal services. These systems have inherent uncertainties, and it is essential to account for the customers' behavior towards such uncertainties. The CPT-based analysis reveals that if we add a stage where each customer is sent a private message before she makes her option selection, then we can get improved results. For example, these messages could take the form of selecting a customer at random to receive priority service or discounted

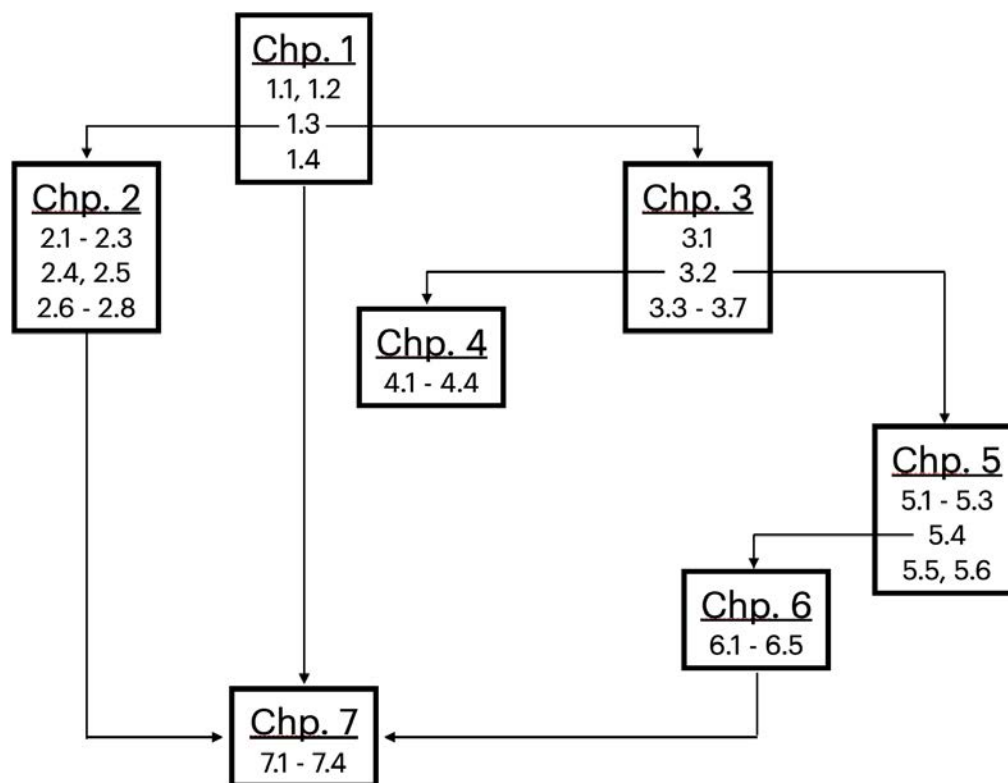


Figure 1.5: Flow chart showing the dependence between different chapters and sections.

pricing. Such messages play the role of nudges that help in aligning the beliefs of the players for optimal service provisioning.

We already observe such nudges and incentives being used around us. But a theory explaining these practices is still in the infant stage. Our mediated mechanism design framework is a very promising direction for explaining these observations theoretically and improving the design of these systems. Mechanism design is commonly referred to as the engineering side of game theory. These theoretical and methodological results can have substantial implications for the design of behavior-aware systems such as online marketplaces and social networks.

We conclude in Chapter 7 with some additional remarks related to the spirit of this work, possible directions for future work, its connections to communication, data analytics, and artificial intelligence, and fairness and ethical considerations related to the use of emotional and psychological aspects in resource allocation.

Appendix 1.A provides the notational conventions followed throughout the document. Figure 1.5 outlines the dependence of the different portions of this thesis on each other.

Appendix

1.A Notational Conventions

We introduce some notational conventions that will be used throughout the rest of the thesis. The scope of any additional notation introduced within a chapter will be limited to that chapter.

Let \mathbb{N} , \mathbb{Z} , \mathbb{Q} , and \mathbb{R} denote the sets of all natural numbers, integers, rational numbers, and real numbers, respectively. Let $\mathbf{1}\{\cdot\}$ denote the indicator function that is equal to one if the predicate inside the brackets is true and is zero otherwise. Let $\text{supp}(\cdot)$ denote the support of the probability distribution within the parentheses. If Z is a subset of a Euclidean space, then let $\text{co}(Z)$ denote the convex hull of Z , and let $\overline{\text{co}}(Z)$ denote the closed convex hull of Z . For any integer $n \in \mathbb{N}$, let $[n] := \{1, 2, \dots, n\}$.

If Z is a Polish space (complete separable metric space), let $\mathcal{P}(Z)$ denote the set of all probability measures on (Z, \mathcal{F}) , where \mathcal{F} is the Borel sigma-algebra of Z . Let $\text{supp}(p)$ denote the support of a distribution $p \in \mathcal{P}(Z)$, i.e. the smallest closed subset of Z such that $p(Z) = 1$. Let $\Delta_f(Z) \subset \mathcal{P}(Z)$ denote the set of all probability distributions that have a finite support. For any element $p \in \Delta_f(Z)$, we will denote the probability of $z \in Z$ assigned under the distribution p by $p(z)$ (or sometimes by $p[z]$). For $z \in Z$, let $\mathbf{1}\{z\} \in \Delta_f(Z)$ denote the probability distribution such that $p(z) = 1$. If Z is finite (and hence a Polish space with respect to the discrete topology), let $\Delta(Z)$ denote the set of all probability distributions on the set Z , viz.

$$\Delta(Z) = \mathcal{P}(Z) = \Delta_f(Z) = \left\{ (p(z))_{z \in Z} \mid p(z) \geq 0 \forall z \in Z, \sum_{z \in Z} p(z) = 1 \right\},$$

with the usual topology. Let Δ^{m-1} denote the standard $(m-1)$ -simplex, i.e. $\Delta([m])$. For a function $f : X \rightarrow \Delta(Y)$, let $f(y|x) = f(x)(y)$ denote the probability of y under the probability distribution $f(x)$.

Let

$$L = \{(p_1, z_1); \dots; (p_t, z_t)\}.$$

denote a lottery with outcomes z_j , $1 \leq j \leq t$, with their corresponding probabilities given by p_j . We assume the lottery to be exhaustive (i.e. $\sum_{j=1}^t p_j = 1$). Note that we are allowed to have $p_j = 0$ for some values of j and we can have $z_k = z_l$ even when $k \neq l$.

If a lottery L consists of a unique outcome z that occurs with probability 1, then with an abuse of notation we will denote the lottery $L = \{(1, z)\}$ simply by $L = z$. Similarly, if a probability distribution $f(x)$ assigns probability 1 to y , then again with an abuse of notation we will write $f(x) = y$. If, for each x , $f(x)$ has a singleton support, then with an abuse of notation we will treat f as a function from X to Y .

Notes

¹Our focus here will be on decision under risk with outcome spaces mapped to real numbers but EUT and CPT extend to more general settings with general outcome sets and subjective beliefs. See [132].

²These assumptions are not needed for the results in this thesis to hold unless stated otherwise.

³ The results in Chapter 2 appear in the paper [109].

⁴ The results in Chapter 3 appear in the paper [108].

⁵ The results in Chapter 4 appear in the pre-print [105].

⁶ The results in Chapter 5 appear in the paper [106].

⁷ The results in Chapter 6 appear in the pre-print [107].

Chapter 2

Optimal Resource Allocation over Networks with CPT Players

2.1 Introduction

We consider the problem of congestion management in a network, and resource allocation amongst heterogeneous users, in particular human agents, with varying preferences. This is a well-recognized problem in network economics [94] and central to most of the examples discussed in Section 1.2. Market-based solutions have proven to be very useful for this purpose, with varied mechanisms, such as auctions and fixed rate pricing [46]. In this chapter, we consider a lottery-based mechanism, as opposed to the deterministic allocations studied in the literature. We mainly ask the following questions: (i) *Do lotteries provide an advantage over deterministic implementations?* (ii) *If yes, then does there exist a market-based mechanism to implement an optimum lottery?*

In order to answer the first question we need to define our goal in allocating resources. There is an extensive literature on the advantages of lotteries: Eckhoff [43] and Stone [125] hold that lotteries are used because of fairness concerns; Boyce [21] argues that lotteries are effective to reduce rent-seeking from speculators; Morgan [87] shows that lotteries are an effective way of financing public goods through voluntary funds, when the entity raising funds lacks tax power; Hylland and Zeckhauser [61] propose implementing lotteries to elicit honest preferences and allocate jobs efficiently. In all of these works, there is an underlying assumption, which is also one of the key reasons for the use of lotteries, that the goods to be allocated are indivisible.

However, we notice lotteries being implemented even when the goods to be allocated are divisible, for example in lottos and parimutuel betting. In several experiments, it has been observed that lottery-based rewards are more appealing than deterministic rewards of the same expected value, and thus provide an advantage in maximizing the desired influence on people's behavior [111]. We also observe several firms presenting lottery-based offers to incentivize customers into buying their products or using their services, and in return to

improve their revenues. Thus, although lottery-based mechanisms are being widely implemented, a theoretical understanding for the same seems to be lacking. This is one of the motivations for this chapter, which aims to justify the use of lottery-based mechanisms, based on models coming from behavioral economics for how humans evaluate options.

2.2 Lottery-Based Resource Allocation Model

We work with the framework proposed by Kelly [72] for throughput control in the Internet with elastic traffic. However, this framework is general enough to have applications to network resource allocation problems arising in several other domains. We have a network with a set $[m] = \{1, \dots, m\}$ of *resources* or *links* and a set $[n] = \{1, \dots, n\}$ of *users* or *players*. Let $c_j > 0$ denote the finite *capacity* of link $j \in [m]$ and let $c := (c_j)_{j \in [m]} \in \mathbb{R}^m$. (All vectors, unless otherwise specified, will be treated as column vectors.) Each user i has a fixed *route* J_i , which is a non-empty subset of $[m]$. Let $R_j := \{i \in [n] | j \in J_i\}$ the set of all players whose route uses link j . We say that an allocation profile x is feasible if it satisfies the capacity constraints of the network, i.e.

$$\sum_{i \in R_j} x_i \leq c_j, \forall j \in [m]. \quad (2.2.1)$$

Let \mathcal{F} denote the set of all feasible allocation profiles. We assume that the network constraints are such that \mathcal{F} is bounded, and hence a polytope.

Instead of allocating a fixed throughput x_i to player $i \in [n]$, we consider allocating her a *lottery* (or a *prospect*)

$$L_i := \{(p_i(1), y_i(1)), \dots, (p_i(k_i), y_i(k_i))\}, \quad (2.2.2)$$

where $y_i(l_i) \geq 0, l_i \in [k_i]$, denotes a throughput and $p_i(l_i), l_i \in [k_i]$, is the probability with which throughput $y_i(l_i)$ is allocated.

We now describe the CPT model we use to measure the “utility” or “happiness” derived by each player from her lottery. In order to focus on the effects of probabilistic sensitivity, and to avoid the complications resulting from reference point considerations, we assume that the reference point of all the players is equal to 0, and we consider prospects with only nonnegative outcomes. This is, in fact, identical to the rank dependent utility (RDU) model [114]. As a result, we assume that each player i is associated with a value function $v_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ that is continuous, differentiable, concave, and strictly increasing, and a probability weighting function $w_i : [0, 1] \rightarrow [0, 1]$ that is continuous, strictly increasing and satisfies $w_i(0) = 0$ and $w_i(1) = 1$.⁸

For the prospect L_i in (2.2.2), let $\pi_i : [k_i] \rightarrow [k_i]$ be a permutation such that

$$z_i(1) \geq z_i(2) \geq \dots \geq z_i(k_i), \quad (2.2.3)$$

and

$$y_i(l_i) = z_i(\pi_i(l_i)) \text{ for all } l_i \in [k_i]. \quad (2.2.4)$$

The prospect L_i can equivalently be written as

$$L_i = \{(\tilde{p}_i(1), z_i(1)); \dots; (\tilde{p}_i(k_i), z_i(k_i))\},$$

where $\tilde{p}_i(l_i) := p_i(\pi_i^{-1}(l_i))$ for all $l_i \in [k_i]$. The *CPT value* of prospect L_i for player i is evaluated using the value function $v_i(\cdot)$ and the probability weighting function $w_i(\cdot)$ as follows:

$$V_i(L_i) = \sum_{l_i=1}^{k_i} d_{l_i}(p_i, \pi_i) v_i(z_i(l_i)), \quad (2.2.5)$$

where $d_{l_i}(p_i, \pi_i)$ are the *decision weights* given by $d_1(p_i, \pi_i) := w_i(\tilde{p}_i(1))$ and

$$d_{l_i}(p_i, \pi_i) := w_i(\tilde{p}_i(1) + \dots + \tilde{p}_i(l_i)) - w_i(\tilde{p}_i(1) + \dots + \tilde{p}_i(l_i - 1)),$$

for $1 < l_i \leq k_i$. And equivalently, the CPT value of prospect L_i , can be written as

$$V_i(L_i) = \sum_{l_i=1}^{k_i} w_i\left(\sum_{s_i=1}^{l_i} \tilde{p}_i(s_i)\right) [v_i(z_i(l_i)) - v_i(z_i(l_i + 1))],$$

where $z_i(k_i + 1) := 0$. Thus the lowest allocation $z_i(k_i)$ is weighted by $w_i(1) = 1$, and every increment in the value of the allocations, $v_i(z_i(l_i)) - v_i(z_i(l_i + 1))$, $\forall l_i \in [k_i - 1]$, is weighted by the probability weighting function of the probability of receiving an allocation at least equal to $z_i(l_i)$.

We take a utilitarian approach of maximizing the ex ante aggregate utility or the net happiness of the players. (See [4] and the references therein for the relation with other goals such as maximizing revenue.) We then ask the question of finding the optimum allocation profile of prospects, one for each user, comprised of throughputs and associated probabilities for that user, that maximizes the aggregate utility of all the players, and is also *feasible*. An allocation profile of prospects for each user is said to be feasible if it can be implemented, i.e. there exists a probability distribution over feasible throughput allocations whose marginals for each player agree with their allocated prospects.

We say that a lottery profile $\{L_1, \dots, L_n\}$ is *feasible* if there exists a joint distribution $p \in \Delta(\prod_i [k_i])$ such that the following conditions are satisfied:

- (i) The marginal distributions agree with L_i for all players i , i.e. $\sum_{l_{-i}} p(l_i, l_{-i}) = p_i(l_i)$ for all $l_i \in [k_i]$, where l_{-i} in the summation ranges over values in $\prod_{i' \neq i} [k_{i'}]$.
- (ii) For each $(l_i)_{i \in [n]} \in \prod_i [k_i]$ in the support of the distribution p (i.e. $p((l_i)_{i \in [n]}) > 0$), the allocation profile $(y_i(l_i))_{i \in [n]}$ is feasible.

Kelly suggested that the throughput allocation problem be posed as one of achieving maximum aggregate utility for the users. In [73], he considers deterministic allocations and each player has a utility function that determines her utility corresponding to an allocation. Its natural analog in our setup can be framed as the following optimization problem:

$$\begin{aligned} & \text{Maximize} && \sum_{i=1}^n V_i(L_i) \\ & \text{subject to} && \text{(i) and (ii)} \end{aligned}$$

The corresponding problem in [73] is a convex optimization problem and permits a decomposition into a central optimization problem and several user optimization problems, one for each user. Based on this decomposition, a market is proposed, in which each user submits an amount she is willing to pay per unit time to the network based on tentative rates that she received from the network; the network accepts these submitted amounts and determines the price of each network link. A user is then allocated a throughput in proportion to her submitted amount and inversely proportional to the sum of the prices of the links she wishes to use. Under certain assumptions, Kelly shows that there exist equilibrium prices and throughput allocations, and that these allocations achieve maximum aggregate utility. Thus the overall *system problem* of maximizing aggregate utility is decomposed into a *network problem* and several *user problems*, one for each individual user. Further, in [73], the authors have proposed two classes of algorithms which can be used to implement a relaxation of the above optimization problem.

The optimization problem in our setting is more complicated than this. The two key reasons for this are the non-convexity of the probability weighting functions and the permutation structure present in the computation of CPT value. In this chapter we will take a closer look at these aspects and see how we can get around these issues. We will obtain a *user-network decomposition* of the optimization problem that would give rise to a market-based mechanism for optimal lottery allocations.

If all the players have EUT utility with concave utility functions, as is typically assumed to model risk-averseness, one can show that there exists a feasible deterministic allocation that achieves the optimum and hence there is no need to consider lotteries. However, if the players' utility is modeled by CPT, then one can improve over the best aggregate utility obtained through deterministic allocations.

For example, Quiggin [115] considers the problem of distributing a fixed amount amongst several homogeneous players with RDU preferences. He concludes that, under certain conditions on players' RDU preferences, the optimum allocation system is a lottery scheme with a few large prizes and a large number of small prizes, and is strictly preferred over distributing the total amount deterministically amongst the players. In Section 2.7, we extend these results to network settings with heterogeneous players.

2.3 Discretization Trick

The distribution p and the throughputs $(y_i(l_i), i \in [n], l_i \in [k_i])$ of a feasible lottery profile together define a *lottery scheme*. In the following, we restrict our attention to specific types of lottery schemes, wherein the network implements with equal probability one of the k

allocation profiles $y(l) := (y_i(l))_{i \in [n]} \in \mathbb{R}_+^n$, for $l \in [k]$. Let $[k] = \{1, \dots, k\}$ denote the set of *outcomes*, where allocation profile $y(l)$ is implemented if outcome l occurs. Clearly, such a scheme is feasible iff each of the allocation profiles $y(l), \forall l \in [k]$ belongs to \mathcal{F} . Player i thus faces the prospect $L_i = \{(1/k, y_i(l))\}_{l=1}^k$ and such a lottery scheme is completely characterized by the tuple $y := (y_i(l), i \in [n], l \in [k])$. By taking k large enough, any lottery scheme can be approximated by such a scheme.

Let $y_i := (y_i(l))_{l \in [k]} \in \mathbb{R}_+^k$. Let $z_i := (z_i(l))_{l \in [k]} \in \mathbb{R}_+^k$ be a vector and $\pi_i : [k] \rightarrow [k]$ be a permutation such that

$$z_i(1) \geq z_i(2) \geq \dots \geq z_i(k),$$

and

$$y_i(l) = z_i(\pi_i(l)) \text{ for all } l \in [k].$$

Note that y_i is completely characterized by π_i and z_i . Then player i 's CPT value will be

$$V_i(L_i) = \sum_{l=1}^k h_i(l) v_i(z_i(l)),$$

where $h_i(l) := w_i(l/k) - w_i((l-1)/k)$ for $l \in [k]$. Let $h_i := (h_i(l))_{l \in [k]} \in \mathbb{R}_+^k$. Note that $h_i(l) > 0$ for all i, l , since the weighting functions are assumed to be strictly increasing.

Looking at the lottery scheme y in terms of individual allocation profiles z_i and permutations π_i for all players $i \in [n]$, allows us to separate those features of y that affect individual preferences and those that pertain to the network implementation. We will later see that the problem of optimizing aggregate utility can be decomposed into two layers: (i) a convex problem that optimizes over resource allocations, and (ii) a non-convex problem that finds the optimal permutation profile.

Let $z := (z_i(l), i \in [n], l \in [k]), \pi := (\pi_i, i \in [n]), h := (h_i(l), i \in [n], l \in [k])$ and $v := (v_i(\cdot), i \in [n])$. Let S_k denote the set of all permutations of $[k]$. The problem of optimizing aggregate utility $\sum_i V_i(L_i)$ subject to the lottery scheme being feasible, can be formulated as follows:

$$\begin{aligned} & \text{SYS}[z, \pi; h, v, A, c] \\ & \text{Maximize} \quad \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\ & \text{subject to} \quad \sum_{i \in R_j} z_i(\pi_i(l)) \leq c_j, \forall j \in [m], \forall l \in [k], \\ & \quad \quad \quad z_i(l) \geq z_i(l+1), \forall i \in [n], \forall l \in [k], \\ & \quad \quad \quad \pi_i \in S_k, \forall i \in [n]. \end{aligned}$$

We set $z_i(k+1) = 0$ for all i , and the $z_i(k+1)$ are not treated as variables. This takes care of the condition $z_i(l) \geq 0$ for all $i \in [n], l \in [k]$.

Note that such discretization serves in two ways:

- Instead of considering the non-convex probability weighting function w_i for each player i , we can restrict our attention to the vector $h_i := (h_i(l), l \in [k])$.
- It highlights the dependence on the permutation structure π .

In addition, the discretization also has a behavioral interpretation. Oftentimes, the players are incapable of discerning the distinction between probabilities that are very close to each other. They also show a poor sense of judgment when it comes to very small probabilities as one in a million. By restricting k say to be 100, we are making sure that the players are faced with lotteries that have integer percentages which they can comprehend better.

Given a permutation profile, the problem of finding optimum feasible throughput allocations is a convex programming problem, which we call the fixed-permutation system problem, and leads to a nice price mechanism. In the next section we prove the existence of equilibrium prices that decompose the fixed-permutation system problem into a network problem and several user problems, one for each player, as in [72]. The prices can be interpreted as the cost imposed on the players and can be implemented in several forms, such as waiting times in waiting-line auctions or first-come-first-served allocations [12, 126], delay or packet loss in the Internet TCP protocol [84, 76], efforts or resources invested by players in a contest [85, 31], or simply money or reward points.

Finding the optimum permutation profile, on the other hand, is a non-convex problem. In Section 2.6, we study the duality gap in the system problem and consider a relaxation of the system problem by allowing the permutations to be doubly stochastic matrices instead of restricting them to be permutation matrices. We show that strong duality holds in the relaxed system problem and so it has value equal to the dual of the original system problem (Theorem 2.6.2). We also consider the problem where link constraints hold in expectation, called the average system problem, and show that strong duality holds in this case and so it has value equal to the relaxed problem. In Section 2.7, we study the average system problem in further detail, and prove a result on the structure of optimal lotteries. We give an example in Section 2.D to show that the duality gap in the original system problem can be nonzero and Theorem 2.6.3 shows that the primal system problem is NP-hard.

2.4 Pricing and Market-Based Mechanism

The system problem $\text{SYS}[z, \pi; h, v, A, c]$ optimizes over z and π . In this section we fix $\pi_i \in S_k$ for all i and optimize over z . Let $\text{SYS_FIX}[z; \pi, h, v, A, c]$ denote this fixed-permutation

system problem.

$$\begin{aligned}
 & \text{SYS_FIX}[z; \pi, h, v, A, c] \\
 & \text{Maximize} \quad \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\
 & \text{subject to} \quad \sum_{i \in R_j} z_i(\pi_i(l)) \leq c_j, \forall j \in [m], \forall l \in [k], \\
 & \quad \quad \quad z_i(l) \geq z_i(l+1), \forall i \in [n], \forall l \in [k].
 \end{aligned}$$

(In contrast with $\text{SYS}(z, \pi; \dots)$, in $\text{SYS_FIX}(z; \pi, \dots)$, the permutation π is thought of as being fixed.) Since $v_i(\cdot)$ is assumed to be a concave function and $h_i(l) > 0$ for all i, l , this problem has a concave objective function with linear constraints. For all $j \in [m], l \in [k]$, let $\lambda_j(l) \geq 0$ be the dual variables corresponding to the constraints $\sum_{i \in R_j} z_i(\pi_i(l)) \leq c_j$ respectively, and for all $i \in [n], l \in [k]$, let $\alpha_i(l) \geq 0$ be the dual variables corresponding to the constraints $z_i(l) \geq z_i(l+1)$ respectively. Let $\lambda := (\lambda_j(l), j \in [m], l \in [k])$ and $\alpha := (\alpha_i(l), i \in [n], l \in [k])$. Then the Lagrangian for the fixed-permutation system problem $\text{SYS_FIX}[z; \pi, h, v, A, c]$ can be written as follows:

$$\begin{aligned}
 \mathcal{L}(z; \alpha, \lambda) &:= \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\
 &+ \sum_{i=1}^n \sum_{l=1}^k \alpha_i(l) [z_i(l) - z_i(l+1)] + \sum_{j=1}^m \sum_{l=1}^k \lambda_j(l) [c_j - \sum_{i \in R_j} z_i(\pi_i(l))] \\
 &= \sum_{i=1}^n \sum_{l=1}^k \left[h_i(l) v_i(z_i(l)) + (\alpha_i(l) - \alpha_i(l-1)) z_i(l) - \left(\sum_{j \in J_i} \lambda_j(\pi_i^{-1}(l)) \right) z_i(l) \right] \\
 &+ \sum_{j=1}^m \sum_{l=1}^k \lambda_j(l) c_j,
 \end{aligned}$$

where $\alpha_i(0) = 0$ for all $i \in [n]$. Differentiating the Lagrangian with respect to $z_i(l)$ we get,

$$\frac{\partial \mathcal{L}(z; \alpha, \lambda)}{\partial z_i(l)} = h_i(l) v'_i(z_i(l)) + \alpha_i(l) - \alpha_i(l-1) - \left(\sum_{j \in J_i} \lambda_j(\pi_i^{-1}(l)) \right).$$

Let

$$\rho_i(l) := \sum_{j \in J_i} \lambda_j(\pi_i^{-1}(l)), \tag{2.4.1}$$

for all $i \in [n], l \in [k]$. This can be interpreted as the price per unit throughput for player i for her l -th largest allocation $z_i(l)$. The price of the lottery z_i for player i is given by

$\sum_{l=1}^k \rho_i(l) z_i(l)$, or equivalently,

$$\sum_{l=1}^k \mu_i(l) [z_i(l) - z_i(l+1)],$$

where

$$\mu_i(l) := \sum_{s=1}^l \rho_i(s), \text{ for all } l \in [k]. \quad (2.4.2)$$

For $l \in [k-1]$, $\alpha_i(l)$ can be interpreted as a transfer of a nonnegative price for player i from her l -th largest allocation to her $(l+1)$ -th largest allocation. Since the allocation $z_i(l+1)$ cannot be greater than the allocation $z_i(l)$, there is a subsidy of $\alpha_i(l)$ in the price of $z_i(l)$ and an equal surcharge of $\alpha_i(l)$ in the price of $z_i(l+1)$. This subsidy and surcharge is nonzero (and hence positive) only if the constraint is binding, i.e. $z_i(l) = z_i(l+1)$. On the other hand, $\alpha_i(k)$ is a subsidy in price given to player i for her lowest allocation, since she cannot be charged anything higher than the marginal utility at her zero allocation.

Let $h_i := (h_i(l))_{l \in [k]} \in \mathbb{R}_+^k$. Consider the following user problem for player i :

$$\begin{aligned} & \text{USER}[m_i; \mu_i, h_i, v_i] \\ & \text{Maximize} \quad \sum_{l=1}^k h_i(l) v_i \left(\sum_{s=l}^k \frac{m_i(s)}{\mu_i(s)} \right) - \sum_{l=1}^k m_i(l) \\ & \text{subject to} \quad m_i(l) \geq 0, \forall l \in [k], \end{aligned} \quad (2.4.3)$$

where $\mu_i := (\mu_i(l), l \in [k])$ is a vector of rates such that

$$0 < \mu_i(1) \leq \mu_i(2) \leq \dots \leq \mu_i(k). \quad (2.4.4)$$

We can interpret this as follows: User i is charged rate $\mu_i(k)$ for her lowest allocation $\delta_i(k) := z_i(k)$. Let $m_i(k)$ denote the budget spent on the lowest allocation and hence $m_i(k) = \mu_i(k) \delta_i(k)$. For $1 \leq l < k$, she is charged rate $\mu_i(l)$ for the additional allocation $\delta_i(l) := z_i(l) - z_i(l+1)$, beyond $z_i(l+1)$ up to the next lowest allocation $z_i(l)$. Let $m_i(l)$ denote the budget spent on l -th additional allocation and hence $m_i(l) = \mu_i(l) \delta_i(l)$.

Let $m := (m_i(l), i \in [n], l \in [k])$ and $\delta := (\delta_i(l), i \in [n], l \in [k])$. Consider the following network problem:

$$\begin{aligned} & \text{NET}[\delta; m, \pi, A, c] \\ & \text{Maximize} \quad \sum_{i=1}^n \sum_{l=1}^k m_i(l) \log(\delta_i(l)) \\ & \text{subject to} \quad \delta_i(l) \geq 0, \forall i, \forall l, \\ & \quad \quad \quad \sum_{i \in R_j} \sum_{s=\pi_i(l)}^k \delta_i(s) \leq c_j, \forall j, \forall l. \end{aligned}$$

This is the well known Eisenberg-Gale convex program [44] and it can be solved efficiently. Kelly et al. [73] proposed continuous time algorithms for finding equilibrium prices and allocations. For results on polynomial time algorithms for these problems see [64, 28]. We have the following decomposition result:

Theorem 2.4.1. *For any fixed π , there exist equilibrium parameters μ^*, m^*, δ^* and z^* such that*

(i) *for each player i , m_i^* solves the user problem $USER[m_i; \mu_i^*, h_i, v_i]$,*

(ii) *δ^* solves the network problem $NET[\delta; m^*, \pi, A, c]$,*

(iii) *$m_i^*(l) = \delta_i^*(l)\mu_i^*(l)$ for all i, l ,*

(iv) *$\delta_i^*(l) = z_i^*(l) - z_i^*(l+1)$ for all i, l , and*

(v) *z^* solves the fixed-permutation system problem $SYS_FIX[z; \pi, h, v, A, c]$.*

The proof of this can be found in Appendix 2.A.

Thus the fixed-permutation system problem can be decomposed into user problems – one for each player – and a network problem, for any fixed permutation profile π . Similar to the framework in [73], we have an iterative process as follows: The network presents each user i with a rate vector μ_i . Each user solves the user problem $USER[m_i; \mu_i, h_i, v_i]$, and submits their budget vector m_i . The network collects these budget vectors $(m_i)_{i \in [n]}$ and solves the network problem $NET[\delta; m^*, \pi, A, c]$ to get the corresponding allocation z (which can be computed from the incremental allocations δ) and the dual variables λ . The network then computes the rate vectors corresponding to each user from these dual variables as given by (2.4.1) and (2.4.2) and presents it to the users as updated rates. Theorem 2.4.1 shows that the fixed-permutation system problem of maximizing the aggregate utility is solved at the equilibrium of the above iterative process. If the value functions $v_i(\cdot)$ are strictly concave, then one can show that the optimal lottery allocation z^* for the fixed-permutation system problem is unique. However, the dual variables λ , and hence the rates $\mu_i, \forall i$, need not be unique. Nonetheless, if one uses the continuous-time algorithm proposed in [73] to solve the network problem, then a similar analysis as in [73], based on Lyapunov stability, shows that the above iterative process converges to the equilibrium lottery allocation z^* .

One of the permutation profiles, say π^* , solves the system problem. In Section 2.6, we explore this in more detail. However, it is interesting to note that, for any fixed permutation profile π , any deterministic solution is a special case of the lottery scheme y with permutation profile π . Thus, it is guaranteed that the solution of the fixed-permutation system problem for any permutation profile π is at least as good as any deterministic allocation. Here is a simple example, where a lottery-based allocation leads to strict improvement over deterministic allocations.




Figure 2.1: Probability weighting function for the example in Section 2.5. The plot shows the probability weighting function given by equation (1.3.6) (also shown in Section 2.5) with $\gamma = 0.61$. The dotted line shows the identity function for reference.

2.5 A Quick Illustrative Example

Consider a network with n players and a single link with capacity c . Let $n = 10$ and $c = 10$. For all players i , we employ the value functions and weighting functions suggested by Kahneman and Tversky [127], given by

$$v_i(x_i) = x_i^{\beta_i}, \beta_i \in [0, 1],$$

and

$$w_i(p_i) = \frac{p_i^{\gamma_i}}{(p_i^{\gamma_i} + (1-p_i)^{\gamma_i})^{1/\gamma_i}}, \gamma_i \in (0, 1],$$

respectively. We take $\beta_i = 0.88$ and $\gamma_i = 0.61$ for all $i \in [n]$. These parameters were reported as the best fits to the empirical data in [127]. The probability weighting function is displayed in Figure 2.1.

By symmetry and concavity of the value function $v_i(\cdot)$, the optimal deterministic allocation is given by allocating c/n to each player i . The aggregate utility for this allocation is $n * v_1(c/n) = 10$.

Now consider the following lottery allocation: Let $k = n = 10$. Let $\pi_i(l) - 1 = l + i \pmod{k}$ for all $i \in [n]$ and $l \in [k]$. Let $x \in [c/n, c]$ and $z_i(1) = x$ for all $i \in [n]$ and $z_i(l) = (c-x)/(n-1)$ for all $i \in [n]$ and $l = 2, \dots, k$. Note that this is a feasible lottery allocation. Such a lottery scheme can be interpreted as follows: Select a “winning” player uniformly at random from all the players. Allocate her a reward x and equally distribute the remaining reward $c - x$

amongst the rest of the players. The ex ante aggregate utility is given by

$$n * [w_1(1/n)v_1(x) + (1 - w_1(1/n))v_1((c - x)/(n - 1))].$$

This function achieves its maximum equal to 14.1690 at $x = 9.7871$. Thus, the above proposed lottery improves the aggregate utility over any deterministic allocation. The optimum lottery allocation is at least as good as 14.1690.

2.6 Optimum Permutation Profile and Duality Gap

The system problem $\text{SYS}[z, \pi; h, v, A, c]$ can equivalently be formulated as

$$\begin{aligned} \max_{\substack{\pi_i \in S_k \forall i, \\ z: z_i(l) \geq z_i(l+1) \forall i, l}} \min_{\lambda \geq 0} & \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\ & + \sum_{j=1}^m \sum_{l=1}^k \lambda_j(l) \left[c_j - \sum_{i \in R_j} z_i(\pi_i(l)) \right]. \end{aligned} \quad (\text{I})$$

Let W_{ps} denote the value of this problem. It is equal to the optimum value of the system problem $\text{SYS}[z, \pi; h, v, A, c]$. By interchanging the max and min, we obtain the following dual problem:

$$\begin{aligned} \min_{\lambda \geq 0} \max_{\substack{\pi_i \in S_k \forall i, \\ z: z_i(l) \geq z_i(l+1) \forall i, l}} & \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\ & + \sum_{j=1}^m \sum_{l=1}^k \lambda_j(l) \left[c_j - \sum_{i \in R_j} z_i(\pi_i(l)) \right]. \end{aligned} \quad (\text{II})$$

Let W_{ds} denote the value of this dual problem. By weak duality, we know that $W_{ps} \leq W_{ds}$. For a fixed $\lambda \geq 0$ and a fixed z that satisfies $z_i(l) \geq z_i(l+1), \forall i, l$, the optimum permutation profile π in the dual problem (II) should minimize

$$\sum_{j=1}^m \sum_{l=1}^k \lambda_j(l) \sum_{i \in R_j} z_i(\pi_i(l)),$$

which equals

$$\sum_i \sum_l \hat{\rho}_i(l) z_i(\pi_i(l)),$$

Here $\hat{\rho}_i(l) := \sum_{j \in J_i} \lambda_j(l)$, is the price per unit allocation for player i under outcome l . Since the numbers $z_i(l)$ are ordered in descending order, any optimal permutation π_i must satisfy

$$\hat{\rho}_i(\pi_i^{-1}(1)) \leq \hat{\rho}_i(\pi_i^{-1}(2)) \leq \dots \leq \hat{\rho}_i(\pi_i^{-1}(k)). \quad (2.6.1)$$

In other words, any optimal permutation profile π of the dual problem (II) must allocate throughputs in the order opposite to that of the prices $\hat{\rho}_i(l)$.

Lemma 2.6.1. *If strong duality holds between the problems (I) and (II), then any optimum permutation profile π^* satisfies (2.6.1) for all i .*

We prove this lemma in Appendix 2.B. In general, there is a non-zero duality gap between the problems (I) and (II) (see Section 2.D for such an example where the optimum permutation profile π^* does not satisfy (2.6.1)).

The permutation π_i can be represented by a $k \times k$ permutation matrix M_i , where $M_i(s, t) = 1$ if $\pi_i(s) = t$ and $M_i(s, t) = 0$ otherwise, for $s, t \in [k]$. The network constraints $\sum_{i \in R_j} z_i(\pi_i(l)) \leq c_j, \forall l \in [k]$, can equivalently be written as $\sum_{i \in R_j} M_i z_i \leq c_j \mathbf{1}$, where $\mathbf{1}$ denotes a vector of appropriate size with all its elements equal to 1, and the inequality is coordinatewise. A possible relaxation of the system problem is to consider doubly stochastic matrices M_i instead of restricting them to be permutation matrices. A matrix is said to be doubly stochastic if all its entries are nonnegative and each row and column sums up to 1. A permutation matrix is hence a doubly stochastic matrix. Let Ω_k denote the set of all doubly stochastic $k \times k$ matrices and let Ω_k^* denote the set of all $k \times k$ permutation matrices.

Let $M = (M_i, i \in [n])$ denote a profile of doubly stochastic matrices. The relaxed system problem can then be written as follows:

$$\begin{aligned}
 & \text{SYS_REL}[z, M; h, v, A, c] \\
 & \text{Maximize} \quad \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\
 & \text{subject to} \quad \sum_{i \in R_j} M_i z_i \leq c_j \mathbf{1}, \forall j, \\
 & \quad \quad \quad z_i(l) \geq z_i(l+1), \forall i, \forall l, \\
 & \quad \quad \quad M_i \in \Omega_k, \forall i.
 \end{aligned}$$

Then the corresponding primal problem can be written as follows:

$$\begin{aligned}
 & \max_{\substack{M_i \in \Omega_k, \forall i, \\ z: z_i(l) \geq z_i(l+1), \forall l, \forall i}} \min_{\lambda_j \geq 0, \forall j} \sum_i \sum_l h_i(l) v_i(z_i(l)) \\
 & \quad \quad \quad + \sum_j \lambda_j^T \left[c_j \mathbf{1} - \sum_{i \in R_j} M_i z_i \right]. \tag{III}
 \end{aligned}$$

where $\lambda_j = (\lambda_j(l))_{l \in [k]} \in \mathbb{R}_+^k$. Let W_{pr} denote the value of this problem. Interchanging min

and max we get the corresponding dual:

$$\begin{aligned} \min_{\lambda_j \geq 0, \forall j} \quad & \max_{\substack{M_i \in \Omega_k \forall i, \\ z: z_i(l) \geq z_i(l+1) \forall l, \forall i}} \sum_i \sum_l h_i(l) v_i(z_i(l)) \\ & + \sum_j \lambda_j^T \left[c_j \mathbf{1} - \sum_{i \in R_j} M_i z_i \right]. \end{aligned} \quad (\text{IV})$$

Let W_{dr} denote the value of this problem. If the link constraints in the relaxed system problem hold then

$$\frac{1}{k} \sum_{i \in R_j} \sum_{l=1}^k z_i(l) = \frac{1}{k} \sum_{i \in R_j} \mathbf{1}^T M_i z_i \leq \frac{1}{k} \mathbf{1}^T c_j \mathbf{1} = c_j. \quad (2.6.2)$$

This inequality essentially says that the link constraints should hold in expectation. Thus we have the following average system problem:

$$\begin{aligned} & \text{SYS_AVG}[z; h, v, A, c] \\ & \text{Maximize} \quad \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\ & \text{subject to} \quad \sum_{i \in R_j} \frac{1}{k} \sum_{l=1}^k z_i(l) \leq c_j, \forall j, \\ & \quad \quad \quad z_i(l) \geq z_i(l+1), \forall i, \forall l, \end{aligned}$$

with its corresponding primal problem:

$$\begin{aligned} \max_{z: z_i(l) \geq z_i(l+1) \forall l, \forall i} \quad & \min_{\bar{\lambda}_j \geq 0, \forall j} \sum_i \sum_l h_i(l) v_i(z_i(l)) \\ & + \sum_j \bar{\lambda}_j \left[c_j - \sum_{i \in R_j} \frac{1}{k} \sum_{l=1}^k z_i(l) \right], \end{aligned} \quad (\text{V})$$

and the dual problem:

$$\begin{aligned} \min_{\bar{\lambda}_j \geq 0, \forall j} \quad & \max_{z: z_i(l) \geq z_i(l+1) \forall l, \forall i} \sum_i \sum_l h_i(l) v_i(z_i(l)) \\ & + \sum_j \bar{\lambda}_j \left[c_j - \sum_{i \in R_j} \frac{1}{k} \sum_{l=1}^k z_i(l) \right], \end{aligned} \quad (\text{VI})$$

where $\bar{\lambda}_j \in \mathbb{R}$ are the dual variables corresponding to the link constraints. Let W_{pa} and W_{da} denote the values of these primal and dual problems respectively.

Then we have the following relation:

Theorem 2.6.2. *For any system problem defined by h, v, A and c , we have*

$$W_{ps} \leq W_{pr} = W_{pa} = W_{da} = W_{dr} = W_{ds}.$$

Proof given in Appendix 2.C.

Thus the duality gap is a manifestation of the “hard” link constraints. In the proof of the above theorem we saw that the relaxed problem is “equivalent” to the average problem and strong duality holds for this relaxation. We will later study the average problem in further detail (Section 2.7).

We observed earlier in Lemma 2.6.1 that if strong duality holds in the system problem, then the optimum permutation profile π^* satisfies (2.6.1). Consider a simple example of two players sharing a single link. Suppose that, at the optimum, $\lambda(l)$ are the prices for $l \in [k]$ corresponding to this link under the different outcomes, and suppose not all of these are equal. Then the optimum permutation profile of the dual problem will align both players’ allocations in the same order, i.e. the high allocations of player 1 will be aligned with the high allocations of player 2. However, we can directly see from the system problem that an optimum π^* should align the two players’ allocations in opposite order. The example in Appendix 2.D builds on this observation and shows that strong duality need not hold for the system problem.

Finally, we prove the following result in Appendix 2.E.

Theorem 2.6.3. *The primal problem (I) is NP-hard.*

2.7 Average System Problem and Optimal Lottery Structure

Suppose it is enough to ensure that the link constraints are satisfied in expectation, as in the average system problem. As an example, suppose we are allocating resources to the players repeatedly and the links have buffers that allow us to allocate excess resources over these links provided the capacity constraints are satisfied on average. If the preferences of the players are not changing with time in this repeated setting, then we get the average system problem. More generally, we should allow the players’ preferences (and perhaps also the capacity constraints) to change with time. We do not consider such a general setting here, however, as the average system problem solves a special case of this problem it would be helpful towards solving the general problem.

Consider the function $V_i^{\text{avg}}(\bar{z}_i)$ on \mathbb{R}_+ given by the value of the following optimization

problem:

$$\begin{aligned}
& \text{Maximize} && \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\
& \text{subject to} && \frac{1}{k} \sum_{l=1}^k z_i(l) = \bar{z}_i, \\
& && z_i(l) \geq z_i(l+1), \forall l \in [k].
\end{aligned} \tag{VII}$$

Let $Z_i(\bar{z}_i)$ denote the set of feasible $(z_i(l))_{l \in [k]}$ in the above problem for any fixed $\bar{z}_i \geq 0$. We observe that $Z_i(\bar{z}_i)$ is a closed and bounded polytope, and hence $V_i^{\text{avg}}(\bar{z}_i)$ is well defined.

Lemma 2.7.1. *For any continuous, differentiable, concave and strictly increasing value function $v_i(\cdot)$, the function $V_i^{\text{avg}}(\cdot)$ is continuous, differentiable, concave and strictly increasing in \bar{z}_i .*

We prove this lemma in Appendix 2.F. The average system problem $\text{SYS_AVG}[z; h, v, A, c]$ can be written as

$$\begin{aligned}
& \text{Maximize} && \sum_{i=1}^n V_i^{\text{avg}}(\bar{z}_i) \\
& \text{subject to} && \sum_{i \in R_j} \bar{z}_i \leq c_j, \forall j, \\
& && \bar{z}_i \geq 0, \forall i.
\end{aligned}$$

Kelly [72] showed that this problem can be decomposed into user problems, one for each user i ,

$$\begin{aligned}
& \text{Maximize} && V_i^{\text{avg}}(\bar{z}_i) - \bar{\rho}_i \bar{z}_i \\
& \text{subject to} && \bar{z}_i \geq 0,
\end{aligned}$$

and a network problem,

$$\begin{aligned}
& \text{Maximize} && \sum_{i=1}^n \bar{\rho}_i \bar{z}_i \\
& \text{subject to} && \sum_{i \in R_j} \bar{z}_i \leq c_j, \forall j, \\
& && \bar{z}_i \geq 0, \forall i,
\end{aligned}$$

in the sense that there exist $\bar{\rho}_i \geq 0, \forall i \in [n]$, such that the optimum solutions \bar{z}_i of the user problems, for each i , solve the network problem and the average system problem. Note that this decomposition is different from the one presented in Section 4.3. Here the network problem aims at maximizing its total revenue $\sum_{i=1}^n \bar{\rho}_i \bar{z}_i$, instead of maximizing a weighted aggregate utility where the utility is replaced with a proxy logarithmic function. The above decomposition is not as useful as the decomposition in Section 4.3 in order to develop iterative

schemes that converge to equilibrium. However, the above decomposition motivates the following user problem:

$$\begin{aligned}
& \text{USER_AVG}[z_i; \bar{\rho}_i, h_i, v_i] \\
& \text{Maximize} && \sum_{l=1}^k h_i(l) v_i(z_i(l)) - \frac{\bar{\rho}_i}{k} \sum_{l=1}^k z_i(l) \\
& \text{subject to} && z_i(l) \geq z_i(l+1), \forall l \in [k],
\end{aligned}$$

where, as before, $z_i(k+1) = 0$.

We observed in Proposition 2.6.2 that strong duality holds in the average system problem. Let z^* be the optimum lottery scheme that solves this problem. Then, first of all, z^* satisfies $z_i^*(l) \geq z_i^*(l+1) \forall i, l$ and is feasible in expectation, i.e., $\bar{z}^* := (\bar{z}_i^*)_{i \in [n]} \in \mathcal{F}$, where $\bar{z}_i^* := (1/k) \sum_l z_i^*(l)$. Further, z^* optimizes the objective function of the average system problem. Besides, there exist $\bar{\lambda}_j^* \geq 0$ for all j such that the primal average problem (V) and the dual average problem (VI) each attain their optimum at z^* , $(\bar{\lambda}_j^*, j \in [m])$.

For player i , consider the price $\bar{\rho}_i^* := \sum_{j \in J_i} \bar{\lambda}_j^*$, which is obtained by summing the prices $\bar{\lambda}_j^*$ corresponding to the links on player i 's route. From the dual average problem (VI), fixing $\bar{\lambda}_j = \bar{\lambda}_j^* \forall j$, we get that the optimum lottery allocation z_i^* for player i should optimize the problem $\text{USER_AVG}[z_i; \bar{\rho}_i^*, h_i, v_i]$.

We now impose some additional conditions on the probability weighting function that are typically assumed based on empirical evidence and certain psychological arguments [67]. We assume that the probability weighting function $w_i(p_i)$ is concave for small values of the probability p_i and convex for the rest. Formally, there exists a probability $\tilde{p}_i \in [0, 1]$ such that $w_i(p_i)$ is concave over the interval $p_i \in [0, \tilde{p}_i]$ and convex over the interval $[\tilde{p}_i, 1]$. Typically the point of inflection, \tilde{p}_i , is around $1/3$.

Let $w_i^* : [0, 1] \rightarrow [0, 1]$ be the minimum concave function that dominates $w_i(\cdot)$, i.e., $w_i^*(p_i) \geq w_i(p_i)$ for all $p_i \in [0, 1]$. Let $p_i^* \in [0, 1]$ be the smallest probability such that $w_i^*(p_i)$ is linear over the interval $[p_i^*, 1]$.

Lemma 2.7.2. *Given the assumptions on $w_i(\cdot)$, we have $p_i^* \leq \tilde{p}_i$ and $w_i^*(p_i) = w_i(p_i)$ for $p_i \in [0, p_i^*]$. If $p_i^* < 1$, then for any $p_i^1 \in [p_i^*, 1)$, we have*

$$w_i(p_i) \leq w_i(p_i^1) + (p_i - p_i^1) \frac{1 - w_i(p_i^1)}{1 - p_i^1}. \quad (2.7.1)$$

for all $p_i \in [p_i^1, 1]$.

A proof of this lemma is included in Appendix 2.G. We now show that, under certain conditions, the optimal lottery allocation z_i^* satisfies

$$z_i^*(l^*) = z_i^*(l^* + 1) = \dots = z_i^*(k), \quad (2.7.2)$$

where $l^* := \min\{l \in [k] : (l-1)/k \geq p_i^*\}$, provided $p_i^* \leq (k-1)/k$. As a result, for a typical optimum lottery allocation, the lowest allocation occurs with a large probability approximately equal to $1 - p_i^*$, and with a few higher allocations that we recognize as bonuses.

Proposition 2.7.3. *For any average user problem $USER_AVG[z_i; \bar{p}_i^*, h_i, v_i]$ with a strictly increasing, continuous, differentiable and strictly concave value function $v_i(\cdot)$, and a strictly increasing continuous probability weighting function $w_i(\cdot)$ (satisfying $w_i(0) = 0$ and $w_i(1) = 1$) such that $p_i^* \leq (k-1)/k$, the optimum lottery allocation z_i^* satisfies Equation (2.7.2).*

Proof of this proposition is provided in Appendix 2.H.

2.8 Summary

We saw that if we take the probabilistic sensitivity of players into account, then lottery allocation improves the ex ante aggregate utility of the players. We considered the RDU model, a special case of CPT utility, to model probabilistic sensitivity. This model, however, is restricted to reward allocations, and it would be interesting to extend it to a general CPT model with reference point and loss aversion. This will allow us to study loss allocations as in punishment or burden allocations, for example criminal justice, military drafting, etc.

For any fixed permutation profile, we showed the existence of equilibrium prices in a market-based mechanism to implement an optimal lottery. We also saw that finding the optimal permutation profile is an NP-hard problem. We note that the system problem has parallels in cross-layer optimization in wireless [79] and multi-route networks [133]. Several heuristic methods have helped achieve approximately optimal solutions in cross-layer optimization. Similar methods need to be developed for our system problem. We leave this for future work.

The hardness in the system problem comes from hard link constraints. Hence, by relaxing these conditions to hold only in expectation, we derived some qualitative features of the optimal lottery structure under the typical assumptions on the probability weighting function of each agent in the RDU model. As observed, the players typically ensure their minimum allocation with high probability, and gamble for higher rewards with low probability.

We assumed that the players are price-takers, i.e. they respond optimally to the prices shown to them in a myopic sense. Such an assumption is reasonable in situations when each player is a small participant in the system and does not have the ability to single-handedly influence the prices. However, more generally, one can imagine the players to behave strategically and could try to manipulate the prices. For example, if a handful of people are competing for a limited resource then they are prone to showing strategic behavior as opposed to the price-taking behavior assumed here. Analyzing such situations requires studying the strategic behavior of the players. In the following chapters, we undertake a systematic study of this by considering games with players having CPT preferences.

Appendix

2.A Proof of Theorem 2.4.1

Since $\text{SYS_FIX}[z; \pi, h, v, A, c]$ is a convex optimization problem, we know that there exist $z^* = (z_i^*(l), i \in [n], l \in [k])$, $\alpha^* = (\alpha_i^*(l), i \in [n], l \in [k])$ and $\lambda^* = (\lambda_j^*(l), j \in [m], l \in [k])$ such that

$$h_i(l)v_i'(z_i^*(l)) = \rho_i^*(l) - \alpha_i^*(l) + \alpha_i^*(l-1), \forall i, \forall l, \quad (2.A.1)$$

$$z_i^*(l) \geq z_i^*(l+1), \quad \alpha_i^*(l) \geq 0, \quad \alpha_i^*(l)(z_i^*(l) - z_i^*(l+1)) = 0, \forall i, \forall l, \quad (2.A.2)$$

$$\sum_{i \in R_j} z_i^*(\pi_i(l)) \leq c_j, \quad \lambda_j^*(l) \geq 0, \quad \lambda_j^*(l)[c_j - \sum_{i \in R_j} z_i^*(\pi_i(l))] = 0, \forall j, \forall l, \quad (2.A.3)$$

where

$$\rho_i^*(l) := \sum_{j \in J_i} \lambda_j^*(\pi_i^{-1}(l)),$$

and such that z^* solves the fixed-permutation system problem $\text{SYS_FIX}[z; \pi, h, v, A, c]$. Hence statement (v) holds for this choice of z^* . Let $\mu_i^*(l) := \sum_{s=1}^l \rho_i^*(s)$, $\delta_i^*(l) := z_i^*(l) - z_i^*(l+1)$ and $m_i^*(l) := \delta_i^*(l)\mu_i^*(l)$ for all i, l . From (2.A.1), we have $\rho_i^*(1) > 0$, because $v_i(\cdot)$ is strictly increasing, $h_i(1) > 0$ and $\alpha_i^*(0) = 0$. Thus, the rate vector μ_i^* satisfies (2.4.4) for all i . Also note that, by construction, the vectors μ^*, δ^*, z^* and m^* satisfy statements (iii) and (iv) of the theorem.

We now show that statement (i) holds. Fix a player i . Observe that the user problem $\text{USER}[m_i; \mu_i^*, h_i, v_i]$ has a concave objective function since $h_i(l) > 0$, $\mu_i^*(l) > 0$ and $v_i(\cdot)$ is concave. Differentiating the objective function of the user problem $\text{USER}[m_i; \mu_i^*, h_i, v_i]$ with respect to $m_i(l)$ at $m_i = m_i^*$, we get

$$\sum_{s=1}^l \frac{h_i(s)v_i'(z_i^*(s))}{\mu_i^*(l)} - 1.$$

From (2.A.1), we have

$$\begin{aligned} \sum_{s=1}^l h_i(s)v_i'(z_i^*(l)) &= \sum_{s=1}^l \rho_i^*(s) - \alpha_i^*(l) \\ &= \mu_i^*(l) - \alpha_i^*(l). \end{aligned}$$

Thus,

$$\sum_{s=1}^l \frac{h_i(s)v_i'(z_i^*(s))}{\mu_i^*(l)} - 1 = -\frac{\alpha_i^*(l)}{\mu_i^*(l)} \leq 0,$$

where equality holds iff $\alpha_i^*(l) = 0$. If $m_i^*(l) > 0$, then $\delta_i^*(l) > 0$ and hence, by (2.A.2), $\alpha_i^*(l) = 0$. Since this holds for all $l \in [k]$, these are precisely the conditions necessary and sufficient for the optimality of m_i^* in the problem of user i , it being a convex problem.

We now show that statement (ii) holds. The Lagrangian corresponding to the network problem $\text{NET}[\delta; m^*, \pi, A, c]$ can be written as follows:

$$\mathcal{L}(\delta; \mu) = \sum_i \sum_{l=1}^k m_i^*(l) \log(\delta_i(l)) + \sum_j \sum_l \mu_j(l) \left[c_j - \sum_{i \in R_j} \sum_{s=\pi_i(l)}^k \delta_i(s) \right],$$

where $\mu_j(l)$ is the dual variable corresponding to the link constraint and $\mu := (\mu_j(l), j \in [n], l \in [k])$. Let $\mu_j(l) = \lambda_j^*(l)$. If $m_i^*(l) > 0$, then $\delta_i^*(l) > 0$, and differentiating the Lagrangian with respect to $\delta_i(l)$ at $\delta_i^*(l)$, we get

$$\begin{aligned} \left. \frac{\partial \mathcal{L}(\delta; \lambda^*)}{\partial \delta_i(l)} \right|_{\delta_i(l)=\delta_i^*(l)} &= \frac{m_i^*(l)}{\delta_i^*(l)} - \sum_{j \in J_i} \sum_{s=1}^l \lambda_j^*(\pi_i^{-1}(s)) \\ &= r_i^*(l) - \sum_{s=1}^l \rho_i^*(s) = 0. \end{aligned}$$

If $m_i^*(l) = 0$, then $\left. \frac{\partial \mathcal{L}(\delta; \lambda^*)}{\partial \delta_i(l)} \right|_{\delta_i(l)=\delta_i^*(l)} \leq 0$ since $\lambda_j^*(l) \geq 0$ for all j, l . Further, from (2.A.3), we have $\lambda_j^*(l) \left[c_j - \sum_{i \in R_j} \sum_{s=\pi_i(l)}^k \delta_i^*(s) \right] = 0$ for all j, l . Thus, δ^* solves the network problem $\text{NET}[\delta; m^*, \pi, A, c]$. \square

2.B Proof of Lemma 2.6.1

Suppose problem (I) and its dual (II) have the same value. The value of (I) is same as that of the system problem $\text{SYS}[z, \pi; h, v, A, c]$. Let us denote the objective function by

$$\begin{aligned} \Theta(\pi, z, \lambda) &:= \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) \\ &\quad + \sum_{j=1}^m \sum_{l=1}^k \lambda_j(l) \left[c_j - \sum_{i \in R_j} z_i(\pi_i(l)) \right]. \end{aligned}$$

Since \mathcal{F} is a polytope, for any fixed permutation profile π , the set of feasible z is closed and bounded. The function $\Theta(\pi, z, \lambda)$ is continuous in z and hence the fixed-permutation system problem $\text{SYS_FIX}[z; \pi, h, v, A, c]$ has a bounded value. We also note that this value is non-negative. Since there are finitely many permutation profiles $\pi \in \prod_i S_k$, maximizing over these, we get that the system problem has a bounded non-negative value, say W , achieved say at z^* and π^* . Thus

$$\sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i^*(l)) = W, \tag{2.B.1}$$

and the lottery z^* is feasible with respect to the permutation profile π^* , i.e. $z_i^*(l) \geq z_i^*(l+1)$ for all $i \in [n], l \in [k]$ and

$$\sum_{i \in R_j} z_i^*(\pi_i^*(l)) \leq c_j \text{ for all } j \in [m], l \in [k]. \quad (2.B.2)$$

If this were not true, then the minimum of the objective function $\Theta(\pi^*, z^*, \lambda)$ with respect to λ would be $-\infty$ and not $W \geq 0$.

The value of the dual problem (II) is equal to W . Consider the function $\Theta_d : \mathbb{R}_+^{m \times k} \rightarrow \mathbb{R}$, given by maximizing over the objective function in problem (II), with respect to π and z for a fixed $\lambda \geq 0$,

$$\Theta_d(\lambda) := \max_{\substack{\pi_i \in S_k \forall i, \\ z: z_i(l) \geq z_i(l+1) \forall i, l}} \Theta(\pi, z, \lambda).$$

We note that the function $\Theta_d(\lambda)$ is lower semi-continuous, since the function $\Theta(\pi, z, \lambda)$ is continuous in λ . Since $(v_i(\cdot), \forall i)$ are concave strictly increasing functions, there exists a sufficiently large finite λ such that $0 \leq M := \Theta_d(\lambda) < \infty$. It follows that the minimum of $\Theta_d(\lambda)$ is achieved over the domain defined by $\lambda_j(l) \in [0, M/(\min_j c_j)]$ for all j, l . Since this is a bounded region and the function $\Theta_d(\lambda)$ is lower semi-continuous, there exists a λ^* such that $\Theta_d(\lambda^*) = \min_{\lambda \geq 0} \Theta_d(\lambda) = W$.

Since $\Theta_d(\lambda^*) = W$, we have $\Theta(\pi^*, z^*, \lambda^*) \leq W$. However, from (2.B.1), (2.B.2) and the fact that $\lambda_j^*(l) \geq 0$ for all j, l we get $\Theta(\pi^*, z^*, \lambda^*) \geq W$. Hence $\Theta(\pi^*, z^*, \lambda^*) = W$. Thus the maximum in the definition of $\Theta_d(\lambda^*)$ is achieved at z^*, π^* . This implies π_i^* satisfies (2.6.1) for all i . \square

2.C Proof of Proposition 2.6.2

As observed earlier, if we replace the condition $M_i \in \Omega_k$ in the primal relaxed problem (III) with the condition $M_i \in \Omega_k^*$, we get the primal system problem (I). Since $\Omega_k^* \subset \Omega_k$, we have $W_{ps} \leq W_{pr}$.

The constraint $\sum_{i \in R_j} \frac{1}{k} \sum_{l=1}^k z_i(l) \leq c_j$ can equivalently be written as

$$\sum_{i \in R_j} \bar{M}_i z_i \leq c_j \mathbf{1},$$

for all j , where \bar{M}_i is the matrix with all its entries equal to $1/k$. Thus, if we replace $M_i \in \Omega_k$ in the primal relaxed problem (III) with \bar{M}_i , we get the primal average problem (V). This implies that $W_{pa} \leq W_{pr}$. However, as observed earlier in Equation (2.6.2), if for some fixed allocations z the link constraints are satisfied with respect to any doubly stochastic matrix M_i , then they are also satisfied with respect to \bar{M}_i . Thus the maximum of the average system problem $\text{SYS_AVG}[z; h, v, A, c]$ is at least as much as the maximum of the relaxed system

problem $\text{SYS_REL}[h, v, A, c]$. Since the maximum of the relaxed system problem is equal to the value of its corresponding primal problem, we get $W_{pa} \geq W_{pr}$. Thus, we have established that $W_{pr} = W_{pa}$.

The average system problem has a concave objective function with linear constraints. Thus, strong duality holds, and we get $W_{pa} = W_{da}$.

We now show that $W_{da} = W_{dr}$. From $W_{pr} \leq W_{dr}$ and $W_{pr} = W_{pa} = W_{da}$, we get $W_{da} \leq W_{dr}$. Suppose $\lambda_j = \bar{\lambda}_j \mathbf{1}$. Then the objective function of the relaxed dual problem (IV),

$$\begin{aligned} & \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) + \sum_{j=1}^m \lambda_j^T \left[c_j \mathbf{1} - \sum_{i \in R_j} M_i z_i \right] \\ &= \sum_{i=1}^n \sum_{l=1}^k h_i(l) v_i(z_i(l)) + \sum_{j=1}^m \bar{\lambda}_j \left[c_j - \sum_{i \in R_j} \frac{1}{k} \sum_{l=1}^k z_i(l) \right], \end{aligned}$$

equals the objective function of the dual average problem (V). This implies that $W_{dr} \leq W_{da}$. This established that $W_{da} = W_{dr}$.

Since any doubly stochastic matrix M_i is a convex combination of permutation matrices by the Birkhoff-von Neumann theorem, in the dual problem (IV), for any fixed λ_j, z_i , the optimum can be achieved by a permutation matrix. This established that $W_{dr} = W_{ds}$.

This completes the proof. \square

2.D Example to Show Duality Gap

Consider the following example with two players $\{1, 2\}$ and a single link with capacity 2.9. Let $k = 2$. Let the corresponding CPT characteristics of the two players be as follows:

$$\begin{aligned} h_1(1) &= \frac{1}{3}, & h_1(2) &= \frac{2}{3}, \\ h_2(1) &= \frac{5}{6}, & h_2(2) &= \frac{1}{6}, \\ v_1(x) &= \log(x + 0.05) + 3, & v_2(x) &= \frac{2 \log(x + 0.05) + 3(x + 0.05)}{5} + 3. \end{aligned}$$

For this problem, it is easy to see that $\pi_1 = (1, 2)$ and $\pi_2 = (2, 1)$ is an optimal permutation. Solving the fixed-permutation system problem with respect to this permutation we get optimal value equal to 7.5621. The corresponding variable values are

$$\begin{aligned} z_1(1) &= y_1(1) = 1.95, & z_1(2) &= y_1(2) = 0.95, \\ z_2(1) &= y_2(2) = 1.95, & z_2(2) &= y_2(1) = 0.95, \end{aligned}$$

and the dual variable values are

$$\lambda_1(1) = \frac{1}{6}, \quad \lambda_1(2) = \frac{2}{3},$$

and $\alpha_i(l) = 0$, for $i = 1, 2, l = 1, 2$. One can check that these satisfy the KKT conditions.

Let us now evaluate the value of the dual problem (II). By symmetry, we can assume without loss of generality that $\lambda_1(1) \leq \lambda_1(2)$. As a result, optimal permutations for the dual problem are given by $\pi_1 = \pi_2 = (1, 2)$. For fixed $\lambda_1(1)$ and $\lambda_1(2)$, we solve the following optimization problem:

$$\begin{aligned}
 \max_{\substack{z_1(1) \geq z_1(2) \geq 0 \\ z_2(1) \geq z_2(2) \geq 0}} & \frac{1}{3} \log(z_1(1) + 0.05) + \frac{2}{3} \log(z_1(2) + 0.05) \\
 & + \frac{5}{6} \left[\frac{2 \log(z_2(1) + 0.05) + 3(z_2(1) + 0.05)}{5} \right] \\
 & + \frac{1}{6} \left[\frac{2 \log(z_2(2) + 0.05) + 3(z_2(2) + 0.05)}{5} \right] \\
 & - \lambda_1(1)[z_1(1) + z_2(1)] - \lambda_1(2)[z_1(2) + z_2(2)] \\
 & + 2.9[\lambda_1(1) + \lambda_1(2)] + 6.
 \end{aligned} \tag{VIII}$$

If $\lambda_1(1) \leq 0.5$, then the value of the problem (VIII) is equal to ∞ (let $z_2(1) \rightarrow \infty$). If $\lambda_1(1) > 0.5$ (and hence $\lambda_1(2) > 0.5$ because $\lambda_1(2) \geq \lambda_1(1)$), then we observe that the effective domain of maximization in the problem (VIII) is compact and problem (VIII) has a finite value. Hence it is enough to consider $\lambda_1(1) > 0.5$. At the optimum there exist $\alpha_1(1), \alpha_1(2), \alpha_2(1), \alpha_2(2) \geq 0$ such that

$$\begin{aligned}
 \lambda_1(1) &= \frac{1}{3} \frac{1}{z_1(1) + 0.05} + \alpha_1(1), \\
 \lambda_1(2) &= \frac{2}{3} \frac{1}{z_1(2) + 0.05} - \alpha_1(1) + \alpha_1(2), \\
 \lambda_1(1) &= \frac{1}{3} \frac{1}{z_2(1) + 0.05} + \frac{1}{2} + \alpha_2(1), \\
 \lambda_1(2) &= \frac{1}{15} \frac{1}{z_2(2) + 0.05} + \frac{1}{10} - \alpha_2(1) + \alpha_2(2),
 \end{aligned}$$

and

$$\begin{aligned}
 \alpha_1(1)[z_1(1) - z_1(2)] &= 0, & \alpha_1(2)z_1(2) &= 0, \\
 \alpha_2(1)[z_2(1) - z_2(2)] &= 0, & \alpha_2(2)z_2(2) &= 0.
 \end{aligned}$$

We now consider each of the sixteen (4×4) cases based on whether the inequalities $z_i(l) \geq z_i(l+1)$ for $i = 1, 2$ and $l = 1, 2$, hold strictly or not.

Case A1 ($z_1(1) = 0, z_1(2) = 0$). Then $\lambda_1(1) \geq 1/0.15$.

Case B1 ($z_1(1) > 0, z_1(2) = 0$). Then $\lambda_1(1) < 1/0.15, \lambda_1(2) \geq 2/0.15$, and

$$\alpha_1(1) = 0, z_1(1) = \frac{1}{3\lambda_1(1)} - 0.05.$$

	B2	D2
C1	10.2284	8.2757
D1	10.1814	9.5006

Table 2.D.1: Table showing the numerical evaluations corresponding to the example in Section 2.D. The numbers in the cells denote the optimum value of the objective function (VIII) in the corresponding cases.

Case C1 ($z_1(1) = z_1(2) > 0$). Then $\lambda_1(2)/2 \leq \lambda_1(1) \leq \lambda_1(2)$, $\lambda_1(1) + \lambda_1(2) < 1/0.05$, and

$$\alpha_1(1) = \frac{2\lambda_1(1) - \lambda_1(2)}{3}, \alpha_1(2) = 0, z_1(1) = z_1(2) = \frac{1}{\lambda_1(1) + \lambda_1(2)} - 0.05.$$

Case D1 ($z_1(1) > z_1(2) > 0$). Then $\lambda_1(1) < \lambda_1(2)/2$, $0 < \lambda_1(1) < 1/0.15$, $0 < \lambda_1(2) < 2/0.15$, and

$$\alpha_1(1) = 0, \alpha_1(2) = 0, z_1(1) = \frac{1}{3\lambda_1(1)} - 0.05, z_1(2) = \frac{2}{3\lambda_1(2)} - 0.05.$$

Case A2 ($z_2(1) = 0, z_2(2) = 0$). Then $\lambda_1(1) \geq (1/0.15) + 0.5$.

Case B2 ($z_2(1) > 0, z_2(2) = 0$). Then $0.5 < \lambda_1(1) < 2.15/0.3$, $\lambda_1(2) \geq (1/0.75) + 0.1$, and

$$\alpha_2(1) = 0, z_2(1) = \frac{2}{6\lambda_1(1) - 3} - 0.05.$$

Case C2 ($z_2(1) = z_2(2) > 0$). This is not possible since $\lambda_1(1) \leq \lambda_1(2)$.

Case D2 ($z_2(1) > z_2(2) > 0$). Then $0.5 < \lambda_1(1) < 2.15/0.3$, $0.1 < \lambda_1(2) < 2.15/1.5$, and

$$\alpha_2(1) = 0, \alpha_2(2) = 0, z_2(1) = \frac{2}{6\lambda_1(1) - 3} - 0.05, z_2(2) = \frac{2}{30\lambda_1(2) - 3} - 0.05.$$

If case A1 or case A2 holds, then $\lambda_1(1) \geq 1/0.15$. For any fixed $\lambda_1(1), \lambda_1(2)$, by choosing $z_i(l), i = 1, 2, l = 1, 2$ small enough (respecting the conditions imposed by the corresponding cases), we get that the value of problem (VIII) is greater than or equal to $2.9 * (1/0.15) = 19.3333 > 7.5621$. Similarly, if case B1 holds, then $\lambda_1(2) \geq 2/0.15$, and we get that the value of problem (VIII) is greater than or equal to $2.9 * (2/0.15) = 38.6667 > 7.5621$. For the remaining 4 cases, substituting the corresponding expressions for $z_i(l), i = 1, 2, l = 1, 2$ in the objective function (VIII) and evaluating the optimum over feasible pairs $(\lambda_1(1), \lambda_1(2))$ for each pair of cases $\{C1, D1\} \times \{B2, D2\}$, the minimum is achieved for the case (C1, D2) and has value equal to 8.2757. Numerical evaluation for each of these cases gives rise to the minimum values as shown in Table 2.D.1. Thus the optimal dual value is 8.2757 and this is strictly greater than the primal value.

□

2.E Proof of Theorem 2.6.3

We describe a polynomial time procedure that reduces an instance of the integer partition problem to a special case of the primal problem. Given a set of positive integers $\{c_1, c_2, \dots, c_n\}$, the integer partition problem is to find a subset $S \subset [n]$, such that

$$\sum_{i \in S} c_i = \sum_{i \notin S} c_i.$$

If such a set S exists, then we say that an integer partition exists. Consider a network with n players and $n + 1$ link constraints given by

$$y_i \leq c_i, \forall i \in [n], \text{ and } \sum_{i=1}^n y_i \leq \frac{\sum_{i=1}^n c_i}{2}.$$

It is easy to realize a network with these link constraints. Let $k = 2$. Let the CPT characteristics of all the players be as follows:

$$h_i(1) = 1 - \epsilon, h_i(2) = \epsilon, v_i(x_i) = x_i, \forall i \in [n],$$

where $\epsilon = 1/10$. Let W_{ps} denote the optimal value of the system problem. We show that $W_{ps} \geq T := (1 - \epsilon) \sum_{i \in [n]} c_i$ if and only if an integer partition exists. Suppose an integer partition exists and is given by the set S , consider the allocation $\pi_i = [1, 2]$ if $i \in S$ and $\pi_i = [2, 1]$ otherwise, $z_i(1) = c_i, z_i(2) = 0$ for all $i \in [n]$. The aggregate utility for this allocation is equal to T and hence $W_{ps} \geq T$. Suppose $W_{ps} \geq T$. Then there an allocation, say z^* and π^* with aggregate utility at least T . Since $k = 2$, π^* actually defines a partition of $[n]$, given by $S = \{i \in [n] : \pi_i(1) = 1\}$. We have, the aggregate utility

$$W(1) + W(2) \geq T,$$

where

$$\begin{aligned} W(1) &:= \sum_{i \in S} (1 - \epsilon) z_i(1) + \sum_{i \notin S} \epsilon z_i(2), \\ W(2) &:= \sum_{i \notin S} (1 - \epsilon) z_i(1) + \sum_{i \in S} \epsilon z_i(2). \end{aligned}$$

Hence at least one of $W(1)$ and $W(2)$ is at least as big as $T/2$. Without loss of generality, let $W(1) \geq T/2$. Thus we have,

$$\sum_{i \in S} z_i(1) + \frac{\epsilon}{1 - \epsilon} \sum_{i \notin S} z_i(2) \geq \frac{\sum_{i \in [n]} c_i}{2}.$$

However, since z^* is feasible, the link constraints give

$$\sum_{i \in S} z_i(1) + \sum_{i \notin S} z_i(2) \leq \frac{\sum_{i \in [n]} c_i}{2}.$$

Since $\epsilon < 1/2$, we should have $\sum_{i \notin S} z_i(2) = 0$ and $\sum_{i \in S} z_i(1) = (\sum_{i \in [n]} c_i)/2$, implying that S forms an integer partition. This completes the proof. \square

2.F Proof of Lemma 2.7.1

Let $\bar{z}_i \geq 0$ and $\tau > 0$. Let $z_i^* \in Z_i(\bar{z}_i)$ be such that $V_i^{\text{avg}}(\bar{z}_i) = \sum_{l=1}^k h_i(l)v_i(z_i^*(l))$. We have, $(z_i^*(l) + \tau)_{l \in [k]} \in Z_i(\bar{z}_i + \tau)$ and

$$V_i^{\text{avg}}(\bar{z}_i + \tau) \geq \sum_{l=1}^k h_i(l)v_i(z_i^*(l) + \tau) > \sum_{l=1}^k h_i(l)v_i(z_i^*(l)) = V_i^{\text{avg}}(\bar{z}_i),$$

where the strict inequality follows from the fact that $v_i(\cdot)$ is strictly increasing. This establishes that the function $V_i^{\text{avg}}(\bar{z}_i)$ is strictly increasing.

Let $\bar{z}_i^1, \bar{z}_i^2 \geq 0$ and $\sigma \in [0, 1]$. Let $z_i^1 \in Z_i(\bar{z}_i^1)$ and $z_i^2 \in Z_i(\bar{z}_i^2)$ be such that $V_i^{\text{avg}}(\bar{z}_i^1) = \sum_{l=1}^k h_i(l)v_i(z_i^1(l))$ and $V_i^{\text{avg}}(\bar{z}_i^2) = \sum_{l=1}^k h_i(l)v_i(z_i^2(l))$. Let $z_i^\sigma := \sigma z_i^1(l) + (1 - \sigma)z_i^2(l)$ and $\bar{z}_i^\sigma := \sigma \bar{z}_i^1 + (1 - \sigma)\bar{z}_i^2$. Then $z_i^\sigma \in Z_i(\bar{z}_i^\sigma)$ and by the concavity of $v_i(\cdot)$, we have

$$\begin{aligned} V_i^{\text{avg}}(\bar{z}_i^\sigma) &\geq \sum_{l=1}^k h_i(l)v_i(z_i^\sigma(l)) \geq \sum_{l=1}^k h_i(l) [\sigma v_i(z_i^1(l)) + (1 - \sigma)v_i(z_i^2(l))] \\ &= \sigma V_i^{\text{avg}}(\bar{z}_i^1) + (1 - \sigma)V_i^{\text{avg}}(\bar{z}_i^2). \end{aligned}$$

This establishes that the function $V_i^{\text{avg}}(\bar{z}_i)$ is concave. This implies that $V_i^{\text{avg}}(\bar{z}_i)$ is continuous and directionally differentiable at each $\bar{z}_i > 0$, and we have the following relation between its left and right directional derivatives (see, for example, [117]):

$$\frac{d}{d\bar{z}_i} V_i^{\text{avg}}(\bar{z}_i-) \geq \frac{d}{d\bar{z}_i} V_i^{\text{avg}}(\bar{z}_i+), \text{ for all } \bar{z}_i > 0. \quad (2.F.1)$$

Further, if $(\bar{z}_i^t)_{t \geq 1}$ is a sequence such that $\bar{z}_i^t \rightarrow 0$, and $z_i^t \in Z_i(\bar{z}_i^t)$ for all $t \geq 1$, then $z_i^t(l) \rightarrow 0$ for all $l \in [k]$. By the continuity of the function $v_i(\cdot)$, we have that the function $V_i^{\text{avg}}(\bar{z}_i)$ is continuous at $\bar{z}_i = 0$.

Let $\bar{z}_i > 0$, and let $\tau^t := 1/t$, for $t \geq 1$. As before, let $z_i^* \in Z_i(\bar{z}_i)$ be such that $V_i^{\text{avg}}(\bar{z}_i) = \sum_{l=1}^k h_i(l)v_i(z_i^*(l))$. Since $\bar{z}_i > 0$ and $(1/k) \sum_{l=1}^k z_i^*(l) = \bar{z}_i$, we have $z_i^*(l) > z_i^*(l+1)$ for at least one $l \in [k]$. Let $\hat{l} \in [k]$ be the smallest such l . For $t \geq 1$, let z_i^{t+} be given by

$$z_i^{t+}(l) := \begin{cases} z_i^*(l) + \frac{k}{\hat{l}} \tau^t, & \text{for } 1 \leq l \leq \hat{l}, \\ z_i^*(l), & \text{for } \hat{l} < l \leq k. \end{cases}$$

Note that $z_i^{t+} \in Z_i(\bar{z}_i + \tau^t)$. We have,

$$\begin{aligned} &V_i^{\text{avg}}(\bar{z}_i + \tau^t) - V_i^{\text{avg}}(\bar{z}_i) - \frac{k\tau^t}{\hat{l}} \sum_{l=1}^{\hat{l}} h_i(l)v_i'(z_i^*(l)) \\ &\geq \sum_{l=1}^k h_i(l)v_i(z_i^{t+}(l)) - \sum_{l=1}^k h_i(l)v_i(z_i^*(l)) - \frac{k\tau^t}{\hat{l}} \sum_{l=1}^{\hat{l}} h_i(l)v_i'(z_i^*(l)) \\ &= \sum_{l=1}^{\hat{l}} h_i(l) \left[v_i(z_i^*(l) + k\tau^t/\hat{l}) - v_i(z_i^*(l)) - \frac{k\tau^t}{\hat{l}} v_i'(z_i^*(l)) \right]. \end{aligned}$$

Let $\gamma^* := \frac{k}{\hat{l}} \sum_{l=1}^{\hat{l}} h_i(l) v_i'(z_i^*(l))$. We have,

$$\frac{d}{d\bar{z}_i} V_i^{\text{avg}}(\bar{z}_i+) \geq \liminf_{t \rightarrow \infty} \frac{V_i^{\text{avg}}(\bar{z}_i + \tau^t) - V_i^{\text{avg}}(\bar{z}_i)}{\tau^t} \geq \gamma^*. \quad (2.F.2)$$

Similarly, for $t \geq \lceil k/z_i^*(\hat{l}) \rceil$ (here, $\lceil \cdot \rceil$ denotes the ceiling function), let z_i^{t-} be given by

$$z_i^{t-}(l) := \begin{cases} z_i^*(l) - \frac{k}{\hat{l}} \tau^t, & \text{for } 1 \leq l \leq \hat{l}, \\ z_i^*(l), & \text{for } \hat{l} < l \leq k. \end{cases}$$

We observe that $z_i^{t-} \in Z_i(\bar{z}_i - \tau^t)$, for all $t \geq \lceil k/z_i^*(\hat{l}) \rceil$, and we have

$$\begin{aligned} & V_i^{\text{avg}}(\bar{z}_i) - V_i^{\text{avg}}(\bar{z}_i - \tau^t) - \frac{k\tau^t}{\hat{l}} \sum_{l=1}^{\hat{l}} h_i(l) v_i'(z_i^*(l)) \\ & \leq \sum_{l=1}^k h_i(l) v_i(z_i^*(l)) - \sum_{l=1}^k h_i(l) v_i(z_i^{t-}(l)) - \frac{k\tau^t}{\hat{l}} \sum_{l=1}^{\hat{l}} h_i(l) v_i'(z_i^*(l)) \\ & = \sum_{l=1}^{\hat{l}} h_i(l) \left[v_i(z_i^*(l)) - v_i(z_i^*(l) - k\tau^t/\hat{l}) - \frac{k\tau^t}{\hat{l}} v_i'(z_i^*(l)) \right]. \end{aligned}$$

This implies,

$$\frac{d}{d\bar{z}_i} V_i^{\text{avg}}(\bar{z}_i-) \leq \limsup_{t \rightarrow \infty} \frac{V_i^{\text{avg}}(\bar{z}_i) - V_i^{\text{avg}}(\bar{z}_i - \tau^t)}{\tau^t} \leq \gamma^*. \quad (2.F.3)$$

From (2.F.1), (2.F.2) and (2.F.3), we have

$$\frac{d}{d\bar{z}_i} V_i^{\text{avg}}(\bar{z}_i-) = \frac{d}{d\bar{z}_i} V_i^{\text{avg}}(\bar{z}_i+) = \gamma^*. \quad (2.F.4)$$

This establishes that the function $V_i^{\text{avg}}(\bar{z}_i)$ is differentiable and completes the proof. \square

2.G Proof of Lemma 2.7.2

Consider the function $\bar{w}_i : [0, 1] \rightarrow [0, 1]$, given by

$$\bar{w}_i(p_i) := \begin{cases} w_i^*(p_i) & \text{for } 0 \leq p_i < \tilde{p}_i, \\ w_i^*(\tilde{p}_i) + (p_i - \tilde{p}_i) \frac{1 - w_i^*(\tilde{p}_i)}{1 - \tilde{p}_i} & \text{for } \tilde{p}_i \leq p_i \leq 1. \end{cases}$$

Since w_i^* is concave on $[0, 1]$, one can verify that the function \bar{w}_i is also concave on $[0, 1]$. Since w_i^* dominates w_i , we have $w_i^*(\tilde{p}_i) \geq w_i(\tilde{p}_i)$. Since the function $w_i(p_i)$ is convex on the interval $[\tilde{p}_i, 1]$, we have $w_i(p_i) \leq \bar{w}_i(p_i)$, for $p_i \in [\tilde{p}_i, 1]$. However, since w_i^* is the minimum

concave function that dominates w_i , we get $\bar{w}_i = w_i^*$. Thus, w_i^* is linear over the interval $[\tilde{p}_i, 1]$, and hence $p_i^* \leq \tilde{p}_i$.

Suppose $\tilde{p}_i = 1$. Then $w_i(\cdot)$ is a concave function on the unit interval $[0, 1]$, and hence $w_i^*(p_i) = w_i(p_i)$, for $p_i \in [0, 1] \supset [0, p_i^*]$. Further, if $p_i^* < 1$, then $w_i(p_i)$ is linear over $[p_i^*, 1]$, and inequality (2.7.1) holds, in fact, with equality. This completes the proof of Lemma 2.7.2, if $\tilde{p}_i = 1$.

For the rest of the proof we assume $\tilde{p}_i < 1$. Define $g_i : [0, 1) \rightarrow \mathbb{R}_+$ as

$$g_i(p_i) := \frac{1 - w_i(p_i)}{1 - p_i}.$$

We now provide an alternate characterization of the function w_i^* and the point p_i^* . Let $\hat{p}_i \in [0, 1]$ be given by

$$\hat{p}_i := \min \arg \min_{p_i \in [0, \tilde{p}_i]} \{g_i(p_i)\}.$$

The existence of \hat{p}_i is guaranteed by the continuity of the function $g_i(p_i)$ on the compact interval $[0, \tilde{p}_i]$. Let $\hat{a}_i := g_i(\hat{p}_i)$. Since $w_i(p_i)$ is convex over the interval $[\tilde{p}_i, 1]$, the function $g_i(p_i)$ is non-decreasing over $[\tilde{p}_i, 1)$. Hence $g_i(p_i) \geq g_i(\hat{p}_i)$, for $p_i \in [0, 1)$. Substituting the expression for $g_i(\cdot)$ and rearranging, we get

$$w_i(p_i) \leq w_i(\hat{p}_i) + \hat{a}_i(p_i - \hat{p}_i),$$

for $p_i \in [0, 1]$. Since the function $w_i(p_i)$ is concave on the interval $[0, \hat{p}_i]$ and the linear function $w_i(\hat{p}_i) + \hat{a}_i(p_i - \hat{p}_i)$ dominates $w_i(p_i)$ on $[0, 1]$, we have that the following function $\hat{w}_i(p_i)$ is concave on $[0, 1]$:

$$\hat{w}_i(p_i) := \begin{cases} w_i(p_i) & \text{for } 0 \leq p_i < \hat{p}_i, \\ w_i(\hat{p}_i) + \hat{a}_i(p_i - \hat{p}_i) & \text{for } \hat{p}_i \leq p_i \leq 1. \end{cases}$$

It follows that $w_i^*(p_i) = \hat{w}_i(p_i)$ for $p_i \in [0, 1]$. Thus, $p_i^* \leq \hat{p}_i$ and $w_i^*(p_i) = \hat{w}_i(p_i) = w_i(p_i)$ for $p_i \in [0, p_i^*]$. If $\hat{p}_i = 0$, then $p_i^* = \hat{p}_i$. If $\hat{p}_i > 0$, then from the definition of \hat{p}_i , we have $g_i(p_i) > g_i(\hat{p}_i)$ for $p_i \in [0, \hat{p}_i)$, and this implies that $\hat{p}_i = p_i^*$.

We now prove inequality (2.7.1). Since $\tilde{p}_i < 1$ we have $p_i^* < 1$. Rearranging we get that inequality (2.7.1) is equivalent to showing that the function $g_i(p_i)$ is non-decreasing over the interval $[p_i^*, 1)$. As observed earlier, $g_i(p_i)$ is non-decreasing over the interval $[\tilde{p}_i, 1)$. Hence it is enough to show that the function $g_i(p_i)$ is non-decreasing on $[p_i^*, \tilde{p}_i]$. Suppose, on the contrary, there exist $p_i^1, p_i^2 \in [p_i^*, \tilde{p}_i]$ such that $p_i^1 < p_i^2$ and $g_i(p_i^1) > g_i(p_i^2)$. Since $p_i^* = \hat{p}_i$, and from the definition of \hat{p}_i , we have $p_i^1 > p_i^*$ and $g_i(p_i^*) \leq g_i(p_i^2)$. Since, $g_i(p_i)$ is a continuous function, there exist $p_i \in [p_i^*, p_i^1]$ such that $g_i(p_i) = g_i(p_i^2)$. Thus we have $p_i < p_i^1 < p_i^2$ such that $g_i(p_i^1) > g_i(p_i) = g_i(p_i^2)$. However, this contradicts the concavity of $w_i(p_i)$ on $[p_i^*, \tilde{p}_i]$. This completes the proof. \square

2.H Proof of Proposition 2.7.3

The Lagrangian for the average user problem $\text{USER_AVG}[z_i; \bar{\rho}_i^*, h_i, v_i]$ is

$$\mathcal{L}(z_i; \alpha_i) = \sum_{l=1}^k h_i(l) v_i(z_i(l)) - \frac{\bar{\rho}_i^*}{k} \sum_{l=1}^k z_i(l) + \sum_{l=1}^k \alpha_i(l) [z_i(l) - z_i(l+1)],$$

where $\alpha_i(l) \geq 0$ are the dual variables corresponding to the order constraints $z_i(l) \geq z_i(l+1)$, and $\alpha_i(0) = 0$. Differentiating with respect to $z_i(l)$, we get,

$$\frac{\partial \mathcal{L}(z_i; \alpha_i)}{\partial z_i(l)} = h_i(l) v_i'(z_i(l)) - \frac{\bar{\rho}_i^*}{k} + \alpha_i(l) - \alpha_i(l-1).$$

Since the problem $\text{USER_AVG}[z_i; \bar{\rho}_i^*, h_i, v_i]$ has a concave objective function and linear constraints, there exist $\alpha_i^*(l) \geq 0$ such that

$$h_i(l) v_i'(z_i^*(l)) = \frac{\bar{\rho}_i^*}{k} - \alpha_i^*(l) + \alpha_i^*(l-1), \forall l \in [k], \quad (2.H.1)$$

and

$$\alpha_i^*(l) [z_i^*(l) - z_i^*(l+1)] = 0, \forall l \in [k]. \quad (2.H.2)$$

If z_i^* consists of identical allocations then it trivially satisfies Equation (2.7.2). If not, then there exists $l^1 \in \{2, \dots, k\}$ such that

$$z_i^*(l^1 - 1) > z_i^*(l^1) = z_i^*(l^1 + 1) = \dots = z_i^*(k),$$

i.e. $z_i^*(l^1)$ is the lowest allocation and occurs with probability $(k - l^1 + 1)/k$, and the next lowest allocation is equal to $z_i^*(l^1 - 1)$. Summing the equations corresponding to $l^1 \leq l \leq k$ from (2.H.1), we get

$$\left[\sum_{s=l^1}^k h_i(s) \right] v_i'(z_i^*(l^1)) = \left(\frac{k - l^1 + 1}{k} \right) \bar{\rho}_i^* - \alpha_i^*(k) + \alpha_i^*(l^1 - 1). \quad (2.H.3)$$

The equation corresponding to $l = l^1 - 1$ in (2.H.1) says,

$$h_i(l^1 - 1) v_i'(z_i^*(l^1 - 1)) = \frac{\bar{\rho}_i^*}{k} - \alpha_i^*(l^1 - 1) + \alpha_i^*(l^1 - 2). \quad (2.H.4)$$

Since $z_i^*(l^1 - 1) > z_i^*(l^1)$, from (2.H.2), we have $\alpha_i^*(l^1 - 1) = 0$. Thus from (2.H.3) and (2.H.4), we have

$$h_i(l^1 - 1) v_i'(z_i^*(l^1 - 1)) \geq \frac{\bar{\rho}_i^*}{k} \geq \frac{1}{k - l^1 + 1} \left[\sum_{s=l^1}^k h_i(s) \right] v_i'(z_i^*(l^1)).$$

Further, since $v_i(\cdot)$ is strictly concave and strictly increasing, $z_i^*(l^1 - 1) > z_i^*(l^1)$ implies $0 < v'_i(z_i^*(l^1 - 1)) < v'_i(z_i^*(l^1))$. Thus,

$$h_i(l^1 - 1) > \frac{1}{k - l^1 + 1} \left[\sum_{s=l^1}^k h_i(s) \right]. \quad (2.H.5)$$

If $(l^1 - 2)/k \geq p_i^*$, then

$$\begin{aligned} h_i(l^1 - 1) &= w_i \left(\frac{l^1 - 1}{k} \right) - w_i \left(\frac{l^1 - 2}{k} \right) \\ &\leq \frac{1}{k - l^1 + 1} \left[w_i(1) - w_i \left(\frac{l^1 - 1}{k} \right) \right] \\ &= \frac{1}{k - l^1 + 1} \left[\sum_{s=l^1}^k h_i(s) \right]. \end{aligned} \quad (2.H.6)$$

where the inequality follows from (2.7.1) with $p_i^1 = (l^1 - 2)/k$ and $p_i = (l^1 - 1)/k$. However, (2.H.6) contradicts (2.H.5) and hence $(l^1 - 2)/k < p_i^*$. This proves the lemma. \square

Notes

⁸ In this chapter, for brevity, we drop the notation of reference point r from v_i^r , and the positive sign from the probability weighting function w_i^+ .

Chapter 3

Notions of Equilibrium: CPT Nash Equilibrium and CPT Correlated Equilibrium

3.1 Introduction

Non-cooperative game theory studies the interaction between decision makers with possibly different objectives. The decision-makers are generally modeled as expected utility maximizers. We will consider games where players have cumulative prospect theoretic preferences. Two of the most well known notions of equilibrium, Nash equilibrium [99] and correlated equilibrium [6], are based on EUT. (See [75] for an excellent account of the strengths and weaknesses of these notions.) Keskin [74] defines analogs for both these equilibrium notions based on cumulative prospect theory. He also establishes the existence of such equilibria under certain continuity conditions. In this chapter, we further study several interesting properties of these notions of equilibria.

There has been considerable interest in the study of the comparative geometry of Nash and correlated equilibria. Under EUT, it is known that the set of all correlated equilibria is a convex polytope and contains the set of all Nash equilibria. In the paper [101], it has been proved that:

the Nash equilibria all lie on the boundary of the correlated equilibrium polytope. (P)

Further, it has been found that in 2-player (bimatrix) games, all extremal Nash equilibria are also extremal correlated equilibria [38, 26, 45], although this result does not hold for more than 2 players [101]. We give a complete characterization of the sets of correlated and Nash equilibria for 2x2 games under CPT, with EUT being a special case.⁹

CPT is known to share common features with EUT. Indeed, recall that CPT is a generalization of EUT. The purpose of this chapter is to study how the geometry of equilibrium notions is affected by prospect theoretic preferences. For example, under CPT, it continues

to be the case that the set of correlated equilibria contains all Nash equilibria, but the set of correlated equilibria is not guaranteed to be a convex polytope (see Example 2 in [74]). The pure Nash equilibria, if they exist, coincide under EUT and CPT (see Proposition 2 in [74]). It is known that the set of correlated equilibria under CPT includes the set of joint probability distributions induced by the convex hull of the set of pure Nash equilibria (see Proposition 3 in [74]), as is true under EUT.

These similarities and differences raise the natural question of whether property (P) continues to hold or not under CPT. In fact, we will see that the set of correlated equilibria can be disconnected (Section 3.6). Nevertheless, our main result says that property (P) continues to hold under CPT (Section 3.3). We also show that for 2×2 games the set of correlated equilibria under CPT is a convex polytope, and we characterize it (Section 3.4).

3.2 Definitions

Let $\Gamma = ([n], (A_i)_{i \in N}, (x_i)_{i \in N})$ be a finite n -person normal form game, where $[n] = \{1, \dots, n\}$ is the set of *players*, A_i is the finite action set of player $i \in [n]$, and $x_i : A_1 \times \dots \times A_n \rightarrow \mathbb{R}$ is the payoff function for player $i \in [n]$. Let each player $i \in [n]$ have at least two actions, i.e. $|A_i| \geq 2, \forall i \in [n]$. Let the set of all action profiles be denoted by $A = \prod_{i \in [n]} A_i$. Let $a_i \in A_i$ denote an action of player $i \in [n]$ (also referred to as a pure strategy or simply a strategy) and let $a = (a_1, \dots, a_n) \in A$ denote an action profile of all players. Let $A_{-i} = \prod_{j \neq i} A_j$ denote the set of action profiles $a_{-i} \in A_{-i}$ of all players except player i . Let $x_i(a)$ denote the payoff of player i when action profile a is played, and let $x_i(\tilde{a}_i, a_{-i})$ denote the payoff to player i when she chooses action $\tilde{a}_i \in A_i$ while the others adhere to a_{-i} .

Definition 3.2.1. The game Γ is *non-trivial* if $x_i(a) \neq x_i(\tilde{a}_i, a_{-i})$ for some player $i \in [n]$, some $a_{-i} \in A_{-i}$, and some $a_i, \tilde{a}_i \in A_i$.

Definition 3.2.2 ([5]). A joint probability distribution $\mu \in \Delta^{|A|-1}$ is said to be a *correlated equilibrium* of Γ if it satisfies the following inequalities:

$$\sum_{a_{-i} \in A_{-i}} \mu(a) (x_i(a_i, a_{-i}) - x_i(\tilde{a}_i, a_{-i})) \geq 0, \text{ for all } i \text{ and for all } a_i, \tilde{a}_i \in A_i. \quad (3.2.1)$$

The set of all correlated equilibria, henceforth denoted as $C_{EUT}(\Gamma)$, is a convex polytope which is a proper subset of $\Delta^{|A|-1}$ iff the game is non-trivial. The set $\Delta^*(A)$ of all joint probability distributions that are of product form is defined by a system of nonlinear constraints, viz.

$$\Delta^*(A) := \{\mu \in \Delta^{|A|-1} : \mu(a) = \mu_1(a_1) \times \dots \times \mu_n(a_n) \quad \forall a \in A\}, \quad (3.2.2)$$

where μ_i denotes the marginal probability distribution on A_i induced by μ . The set of all Nash equilibria is the intersection of $\Delta^*(A)$ and $C_{EUT}(\Gamma)$, which is non-empty by virtue of Nash's existence theorem.

We now describe the notion of correlated equilibrium incorporating CPT preferences, as defined by Keskin [74]. Let $\{v_i^{r_i}(\cdot), r_i \in \mathbb{R}\}$ be a family of value functions, one for each

reference point, and $w_i^\pm(\cdot)$ be the probability weighting functions for each player $i \in [n]$. We assume that $v_i^{r_i}(z)$ is continuous in z and r_i for each i . For every player $i \in N$, let the reference point be determined by a continuous function $r_i : \Delta^{|A|-1} \rightarrow \mathbb{R}$. Let $V_i^{r_i}(L)$ denote the CPT value of a lottery L evaluated by player i , using the value function $v_i^{r_i}(\cdot)$ and probability weighting functions $w_i^\pm(\cdot)$ as described in equation (1.3.8) or (1.3.9).

Corresponding to a lottery

$$L = \{(p_1, z_1); \dots; (p_t, z_t)\},$$

as in equation (1.3.1), let $z := (z_1, \dots, z_t)$ and $p := (p_1, \dots, p_t)$. We denote L as (p, z) and refer to the vector z as an *outcome profile*.

For a joint distribution $\mu \in \Delta^{|A|-1}$, let

$$\mu_i(a_i) = \sum_{a_{-i} \in A_{-i}} \mu(a_i, a_{-i})$$

be the marginal distribution of player i , and for a_i such that $\mu_i(a_i) > 0$ let

$$\mu_{-i}(a_{-i}|a_i) = \frac{\mu(a_i, a_{-i})}{\mu_i(a_i)}$$

be the conditional distribution on A_{-i} .

If player i observes a signal to play a_i drawn from the joint distribution μ , and if she decides to deviate to an action $\tilde{a}_i \in A_i$, then she will face the lottery

$$L(\mu, a_i, \tilde{a}_i) := \{(\mu_{-i}(a_{-i}|a_i), x_i(\tilde{a}_i, a_{-i}))\}_{a_{-i} \in A_{-i}}.$$

Definition 3.2.3 ([74]). A joint probability distribution $\mu \in \Delta^{|A|-1}$ is said to be a *CPT correlated equilibrium* of Γ if it satisfies the following inequalities for all i and for all $a_i, \tilde{a}_i \in A_i$ such that $\mu_i(a_i) > 0$:

$$V_i^{r_i(\mu)}(L(\mu, a_i, a_i)) \geq V_i^{r_i(\mu)}(L(\mu, a_i, \tilde{a}_i)). \quad (3.2.3)$$

Let $C(\Gamma)$ denote the set of all CPT correlated equilibria of Γ .

For any fixed reference point r , since the value function $v^r(\cdot)$ is assumed to be strictly increasing, one can check that two outcome profiles z and y have equal CPT value under all probability distributions p , i.e. $V^r(p, z) = V^r(p, y)$ for all p , iff $z = y$. It then follows that the set $C(\Gamma)$ is a proper subset of $\Delta^{|A|-1}$ iff the game is non-trivial.

We now describe the notion of CPT Nash equilibrium as defined by Keskin¹⁰ [74]. For a mixed strategy $\mu \in \Delta^*(A)$, if player i decides to play a_i , drawn from the distribution μ_i , then she will face the lottery

$$L(\mu_{-i}, a_i) := \{(\mu_{-i}(a_{-i}), x_i(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}},$$

where $\mu_{-i}(a_{-i}) = \prod_{j \neq i} \mu_j(a_j)$ plays the role of $\mu_{-i}(a_{-i}|a_i)$, which does not depend on a_i . Suppose player i decides to deviate and play a mixed strategy μ'_i while the rest of the players continue to play μ_{-i} . Then define the average CPT value for player i by

$$\mathcal{A}_i(\mu'_i, \mu_{-i}) = \sum_{a_i \in A_i} \mu'_i(a_i) V_i^{r_i(\mu)}(L(\mu_{-i}, a_i)).$$

The best response of player i to a mixed strategy $\mu \in \Delta^*(A)$ is defined as

$$BR_i(\mu) := \{\mu_i^* \in \Delta^{|A_i|-1} \mid \forall \mu'_i \in \Delta^{|A_i|-1}, \mathcal{A}_i(\mu_i^*, \mu_{-i}) \geq \mathcal{A}_i(\mu'_i, \mu_{-i})\}. \quad (3.2.4)$$

Definition 3.2.4 ([74]). A mixed strategy $\mu^* \in \Delta^*(A)$, is a CPT Nash equilibrium iff

$$\mu_i^* \in BR_i(\mu^*) \text{ for all } i.$$

We call μ^* a pure or mixed CPT Nash equilibrium depending on μ^* being a pure or mixed strategy respectively.

The set of all CPT correlated equilibria is no longer guaranteed to be a convex polytope (Example 2 in [74]). The set of all CPT Nash equilibria is the intersection of $\Delta^*(A)$ and $C(\Gamma)$ (Proposition 1 in [74]) and is non-empty (Theorem 1 in [74]). We are interested in studying the geometry of this intersection. It should be noted that the set $C(\Gamma)$ depends on the choice of the reference functions $r_i(\mu)$, as does the set of CPT Nash equilibria.

3.3 Main Result: An Interesting Geometric Property

In the case of traditional utility-theoretic equilibria, it has been proved that

Proposition 3.3.1 ([101]). *In any finite, non-trivial game, the Nash equilibria are on the boundary of the polytope of correlated equilibria when it is viewed as a subset of the smallest affine space containing all joint probability distributions.*

Since the set of correlated equilibria $C_{EUT}(\Gamma)$ is a convex polytope, it is enough to prove that the Nash equilibria lie on one of the faces of $C_{EUT}(\Gamma)$ if $C_{EUT}(\Gamma)$ is full-dimensional, i.e. has dimension $|A| - 1$, when it is viewed as a subset of the affine space containing $\Delta^{|A|-1}$, and the statement is trivially true if it is not full-dimensional. When the set $C_{EUT}(\Gamma)$ is not full-dimensional, it is possible for the Nash equilibria to lie in the relative interior of the set $C_{EUT}(\Gamma)$ (Proposition 2 in [101]). Further, the class of games with the Nash equilibrium in the relative interior of the correlated equilibrium polytope has been characterized in the paper [128].

We now extend the above proposition for equilibria with CPT preferences. The proof is quite different since in general $C(\Gamma)$ is not a convex polytope, as shown in Section 3.6 below (see also Example 2 in [74]).

Proposition 3.3.2. *In any finite, non-trivial game, the CPT Nash equilibria are on the boundary of the set of CPT correlated equilibria set when it is viewed as a subset of the smallest affine space containing all joint probability distributions.*

We first prove a lemma which in itself is an interesting property of cumulative prospect theoretic preferences. Let $V^r(\cdot)$ denote the CPT value evaluated with respect to a value function $v^r(\cdot)$ and probability weighting functions $w^\pm(\cdot)$ with respect to a reference point $r \in \mathbb{R}$. Let $z = (z_1, \dots, z_t)$ and $y = (y_1, \dots, y_t)$ be two outcome profiles and $p = (p_1, \dots, p_t)$ be a probability distribution. The prospect (p, z) is said to *pointwise dominate* the prospect (p, y) if $z_j \geq y_j$ for all j such that $p_j > 0$. Further, if the inequality $z_j \geq y_j$ holds strictly for at least one j with $p_j > 0$ then the prospect (p, z) is said to *strictly pointwise dominate* the prospect (p, y) . Let the *regret* corresponding to choosing y instead of z be denoted by

$$\mathcal{R}^r(p, z, y) := V^r(p, z) - V^r(p, y). \quad (3.3.1)$$

Prospects (p, z) and (p, y) are said to be *similarly ranked* if there exists a permutation $(\alpha_1, \dots, \alpha_t)$ of $T' := \{j \in \{1, \dots, t\} | p_j > 0\}$ such that

$$z_{\alpha_1} \geq \dots \geq z_{\alpha_t} \text{ and } y_{\alpha_1} \geq \dots \geq y_{\alpha_t}.$$

Lemma 3.3.3. *In the above setting, suppose the prospects (p, z) and (p, y) satisfy either of the following:*

- (i) *they are not similarly ranked or,*
- (ii) *neither of them dominates the other,*

then there exists a direction $\delta = (\delta_1, \dots, \delta_t)$ with $\sum_{j=1}^t \delta_j = 0$ and $\delta_j = 0$ for $j \notin T'$ such that

$$\mathcal{R}^r(p + \epsilon\delta, z, y) < \mathcal{R}^r(p, z, y) \quad (3.3.2)$$

for all $r \in \mathbb{R}$, for all $\epsilon > 0$ such that $p + \epsilon\delta \in \Delta^{t-1}$.

Proof. We observe that it is enough to prove the claim for the case when $p_j > 0$ for all $1 \leq j \leq t$ because if not, then we can let z', y' and p' be respectively the vectors z, y and p restricted to the coordinates in T' and then use the result. WLOG let the ordering be such that $z_1 \geq \dots \geq z_t$. Let $\delta(j_1, j_2)$ correspond to transferring probability from j_1 to j_2 , i.e. for all $1 \leq j \leq t$,

$$\delta_j(j_1, j_2) = \begin{cases} 1 & \text{if } j = j_2, \\ -1 & \text{if } j = j_1, \\ 0 & \text{otherwise.} \end{cases}$$

Suppose (i) holds. Then there exists $j_1 < j_2$ (and hence $z_{j_1} \geq z_{j_2}$) such that $y_{j_1} < y_{j_2}$. Now, by the strict stochastic dominance property of CPT we have

$$V^r(p + \epsilon\delta, z) \leq V^r(p, z) \text{ and } V^r(p, y) < V^r(p + \epsilon\delta, y),$$

where δ denotes $\delta(j_1, j_2)$, and hence (3.3.2) follows.

Now suppose (p, z) and (p, y) are similarly ranked. WLOG let the ordering be such that $z_1 \geq \dots \geq z_t$ and $y_1 \geq \dots \geq y_t$. Suppose (ii) holds. Then there exist j_1, j_2 such that $z_{j_1} > y_{j_1}$ and $z_{j_2} < y_{j_2}$. In fact, one can find j_1, j_2 such that $z_{j_1} > y_{j_1}, z_j = y_j$ for all j between j_1 and j_2 , and $z_{j_2} < y_{j_2}$. Depending on the order of j_1 and j_2 we have the following two cases (note $j_1 \neq j_2$):

Case 1 ($j_1 < j_2$): Then we have the ordering $z_{j_2} < y_{j_2} \leq y_{j_1} < x_{j_1}$. Let $\delta = \delta(j_1, j_2)$. Then it follows from the strict monotonicity of the functions $w_i^\pm(\cdot)$ and the definition of decision weights that

$$\begin{aligned} \nabla_{j_1}^+(p + \epsilon\delta) - \nabla_{j_1}^+(p) &< 0, \\ \nabla_{j_2}^+(p + \epsilon\delta) - \nabla_{j_2}^+(p) &> 0, \\ \nabla_{j_1}^-(p + \epsilon\delta) - \nabla_{j_1}^-(p) &< 0, \\ \nabla_{j_2}^-(p + \epsilon\delta) - \nabla_{j_2}^-(p) &> 0. \end{aligned}$$

(We suppress the dependence of $\nabla_j^\pm(p, a)$ on the permutation a since we have assumed z and y to be ordered.) Depending on the position of the reference point r , we have the following subcases:

Subcase 1a ($r \leq x_{j_2}$):

$$\begin{aligned} &[V^r(p + \epsilon\delta, x) - V^r(p + \epsilon\delta, y)] - [V^r(p, z) - V^r(p, y)] \\ &= [\nabla_{j_1}^+(p + \epsilon\delta) - \nabla_{j_1}^+(p)][v^r(z_{j_1}) - v^r(y_{j_1})] \\ &\quad + [\nabla_{j_2}^+(p + \epsilon\delta) - \nabla_{j_2}^+(p)][v^r(z_{j_2}) - v^r(y_{j_2})], \end{aligned}$$

because $\nabla_j^+(p + \epsilon\delta) = \nabla_j^+(p)$ for all $j \notin \{j_1, \dots, j_2\}$ and $v^r(z_j) = v^r(y_j)$ for all $j_1 < j < j_2$. Since $v^r(z_{j_1}) - v^r(y_{j_1}) > 0$ and $v^r(z_{j_2}) - v^r(y_{j_2}) < 0$ we get (3.3.2).

Subcase 1b ($x_{j_2} < r \leq y_{j_2}$):

$$\begin{aligned} &[V^r(p + \epsilon\delta, x) - V^r(p + \epsilon\delta, y)] - [V^r(p, x) - V^r(p, y)] \\ &= [\nabla_{j_1}^+(p + \epsilon\delta) - \nabla_{j_1}^+(p)][v^r(z_{j_1}) - v^r(y_{j_1})] + [\nabla_{j_2}^-(p + \epsilon\delta) - \nabla_{j_2}^-(p)]v^r(x_{j_2}) \\ &\quad - [\nabla_{j_2}^+(p + \epsilon\delta) - \nabla_{j_2}^+(p)]v^r(y_{j_2}). \end{aligned}$$

Now $v^r(z_{j_1}) - v^r(y_{j_1}) > 0$, $v^r(z_{j_2}) < 0$, $v^r(y_{j_2}) > 0$ and the result follows.

Subcase 1c ($y_{j_2} < r \leq y_{j_1}$):

$$\begin{aligned} &[V^r(p + \epsilon\delta, x) - V^r(p + \epsilon\delta, y)] - [V^r(p, z) - V^r(p, y)] \\ &= [\nabla_{j_1}^+(p + \epsilon\delta) - \nabla_{j_1}^+(p)][v^r(z_{j_1}) - v^r(y_{j_1})] \\ &\quad + [\nabla_{j_2}^-(p + \epsilon\delta) - \nabla_{j_2}^-(p)][v^r(z_{j_2}) - v^r(y_{j_2})]. \end{aligned}$$

Now $v^r(z_{j_1}) - v^r(y_{j_1}) > 0$, $v^r(z_{j_2}) - v^r(y_{j_2}) < 0$ and the result follows.

Subcase 1d ($y_{j_1} < r \leq z_{j_1}$):

$$\begin{aligned} & [V^r(p + \epsilon\delta, z) - V^r(p + \epsilon\delta, y)] - [V^r(p, z) - V^r(p, y)] \\ &= [\nabla_{j_1}^+(p + \epsilon\delta) - \nabla_{j_1}^+(p)]v^r(z_{j_1}) - [\nabla_{j_1}^-(p + \epsilon\delta) - \nabla_{j_1}^-(p)]v^r(y_{j_1}) \\ & \quad + [\nabla_{j_2}^-(p + \epsilon\delta) - \nabla_{j_2}^-(p)][v^r(z_{j_2}) - v^r(y_{j_2})]. \end{aligned}$$

Now $v^r(z_{j_1}) > 0$, $v^r(y_{j_1}) < 0$, $v^r(z_{j_2}) - v^r(y_{j_2}) < 0$ and the result follows.

Subcase 1e ($x_{j_i} < r$):

$$\begin{aligned} & [V^r(p + \epsilon\delta, x) - V^r(p + \epsilon\delta, y)] - [V^r(p, x) - V^r(p, y)] \\ &= [\nabla_{j_1}^-(p + \epsilon\delta) - \nabla_{j_1}^-(p)][v^r(z_{j_1}) - v^r(y_{j_1})] \\ & \quad + [\nabla_{j_2}^-(p + \epsilon\delta) - \nabla_{j_2}^-(p)][v^r(z_{j_2}) - v^r(y_{j_2})]. \end{aligned}$$

Now $v^r(z_{j_1}) - v^r(y_{j_1}) > 0$, $v^r(z_{j_2}) - v^r(y_{j_2}) < 0$ and the result follows.

Case 2 ($j_1 > j_2$) implies the order $y_{j_1} < x_{j_1} \leq z_{j_2} < y_{j_2}$. Taking $\delta = \delta(j_2, j_1)$, each of the subcases depending on the position of the reference point can be handled as in case 1. \square

Remark 3.3.4. The vector δ used in the proof of this lemma depends only on the prospects (p, z) and (p, y) and not on the reference point r . In fact, it depends only on the order structure of the vectors z and y and not on the probability distribution vector p as long as $p_j > 0$ for all $1 \leq j \leq t$. Also, the range of ϵ for which the claim holds depends only on the prospects (p, z) and (p, y) and not on the reference point r . Lemma 3.3.3 can be extended to more general CPT settings as in [30], where the outcome space is assumed to be a connected topological space instead of monetary outcomes in \mathbb{R} .

Proof of proposition 3.3.2. If a CPT Nash equilibrium $\hat{\mu}$ is not completely mixed, i.e. there is a player i and an action $a_i \in A_i$, such that $\hat{\mu}_i(a_i) = 0$, then $\hat{\mu}$ assigns zero probability to one or more action profiles and hence lies on the boundary of $\Delta^{|A|-1}$ and thus also on the boundary of C_{CPT} .

Suppose now that $\hat{\mu} \in \Delta^*(A) \cap C(\Gamma)$ is completely mixed. Then the inequalities (3.2.3) hold for all i and for all $a_i, \tilde{a}_i \in A_i$. In particular, for any pair $a_i, \tilde{a}_i \in A_i$ we have

$$\begin{aligned} & V_i^{r_i(\hat{\mu})} \left(\{(\hat{\mu}_{-i}(a_{-i}|a_i), x_i(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}} \right) \\ & \geq V_i^{r_i(\hat{\mu})} \left(\{(\hat{\mu}_{-i}(a_{-i}|a_i), x_i(\tilde{a}_i, a_{-i}))\}_{a_{-i} \in A_{-i}} \right), \\ & V_i^{r_i(\hat{\mu})} \left(\{(\hat{\mu}_{-i}(a_{-i}|\tilde{a}_i), x_i(\tilde{a}_i, a_{-i}))\}_{a_{-i} \in A_{-i}} \right) \\ & \geq V_i^{r_i(\hat{\mu})} \left(\{(\hat{\mu}_{-i}(a_{-i}|\tilde{a}_i), x_i(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}} \right). \end{aligned}$$

However, since $\hat{\mu} \in \Delta^*(A)$, we have $\hat{\mu}_{-i} := \hat{\mu}_{-i}(\cdot|a_i) = \hat{\mu}_{-i}(\cdot|\tilde{a}_i)$ and hence the above inequalities are in fact equalities. The same is true for all the inequalities (3.2.3).

Since the game is non-trivial, there exist $i \in [n]$ and $a_i, \tilde{a}_i \in A_i$ such that $x_i(a_i, a_{-i}) \neq x_i(\tilde{a}_i, a_{-i})$ for some $a_{-i} \in A_{-i}$. Consider the inequality in (3.2.3) corresponding to such an

(i, a_i, \tilde{a}_i) . Fix a one to one correspondence between the numbers $\{1, \dots, t\}$ and the action profiles $\{a_{-i} \in A_{-i}\}$ (here $t = |A_{-i}|$). Let

$$z := (z_1, \dots, z_t) = (x_i(a_i, a_{-i}))_{a_{-i} \in A_{-i}},$$

and

$$y := (y_1, \dots, y_t) = (x_i(\tilde{a}_i, a_{-i}))_{a_{-i} \in A_{-i}}.$$

Since $\hat{\mu}$ is completely mixed, $\hat{\mu}_i(a_i) > 0$. Let

$$p = (p_1, \dots, p_t) = (\hat{\mu}_{-i}(a_{-i}))_{a_{-i} \in A_{-i}}$$

be the conditional probability distribution on A_{-i} .

If either profile (p, z) or (p, y) pointwise dominated the other then the pointwise dominance would be strict since z and y are distinct and $p_j > 0$ for all $1 \leq j \leq t$. By the strict monotonicity property of CPT, we would get $V^{r_i(\hat{\mu})}(p, z) \neq V^{r_i(\hat{\mu})}(p, y)$ contrary to our assumption. Thus condition (ii) in Lemma 3.3.3 is satisfied and there exists a direction vector $\delta = (\delta_1, \dots, \delta_t)$ with $\sum_{j=1}^t \delta_j = 0$ such that $V_i^{r_i}(p + \epsilon\delta, z) < V_i^{r_i}(p + \epsilon\delta, y)$ for all $r_i \in \mathbb{R}$, for all $\epsilon > 0$ such that $p + \epsilon\delta \in \Delta^{t-1}$. Note that the vector δ and the range of ϵ does not depend on the reference point r_i (see Remark 3.3.4). Consider the joint probability distribution $\bar{\mu}$ given by

$$\bar{\mu}(\bar{a}_i, a_{-i}) = \begin{cases} \hat{\mu}_i(a_i)(p_j + \epsilon\delta_j) & \text{if } \bar{a}_i = a_i \text{ and } j \text{ corresponds to } a_{-i}, \\ \hat{\mu}(\bar{a}_i, a_{-i}) & \text{otherwise.} \end{cases}$$

Let $\mathcal{R}_i^r(\cdot)$ denote the regret corresponding to player i , evaluated using her value function and probability weighting functions. This should be thought of as defined for any pair of outcome profiles z and y on A_{-i} with a given probability distribution p on A_{-i} , as in equation (3.3.1), with V^r being replaced by V_i^r and defined as in equation (1.3.8), using the value function v_i^r and the weighting functions w_i^\pm . Since $\hat{\mu} \in \Delta^*(A) \cap C(\Gamma)$,

$$\mathcal{R}_i^{r_i(\hat{\mu})}(\hat{\mu}_{-i}, z, y) = V_i^{r_i(\hat{\mu})}(p, z) - V_i^{r_i(\hat{\mu})}(p, y) = 0,$$

and

$$\mathcal{R}_i^{r_i(\hat{\mu})}(\hat{\mu}_{-i}, y, z) = V_i^{r_i(\hat{\mu})}(p, y) - V_i^{r_i(\hat{\mu})}(p, z) = 0.$$

respectively. Now if

$$\mathcal{R}_i^{r_i(\bar{\mu})}(\hat{\mu}_{-i}, z, y) \leq \mathcal{R}_i^{r_i(\hat{\mu})}(\hat{\mu}_{-i}, z, y)$$

then from the choice of $\bar{\mu}$

$$\mathcal{R}_i^{r_i(\bar{\mu})}(\bar{\mu}_{-i}(\cdot|a_i), z, y) = \mathcal{R}_i^{r_i(\bar{\mu})}(p + \epsilon\delta, z, y) < \mathcal{R}_i^{r_i(\bar{\mu})}(p, z, y) = \mathcal{R}_i^{r_i(\bar{\mu})}(\hat{\mu}_{-i}, z, y) \leq 0.$$

On the other hand, if

$$\mathcal{R}_i^{r_i(\bar{\mu})}(\hat{\mu}_{-i}, z, y) > \mathcal{R}_i^{r_i(\hat{\mu})}(\hat{\mu}_{-i}, z, y) = 0$$

then

$$\mathcal{R}_i^{r_i(\bar{\mu})}(\bar{\mu}_{-i}(\cdot|\tilde{a}_i), y, z) = R_i^{r_i(\bar{\mu})}(\hat{\mu}_{-i}, y, z) = -\mathcal{R}_i^{r_i(\bar{\mu})}(\hat{\mu}_{-i}, z, y) < 0.$$

Thus either of the inequalities in (3.2.3) corresponding to deviation from a_i to \tilde{a}_i or \tilde{a}_i to a_i is violated by the joint distribution $\bar{\mu}$. Thus, for any neighborhood N of $\hat{\mu}$, $\bar{\mu}$ belongs to N for sufficiently small ϵ and $\bar{\mu} \notin C(\Gamma)$. Thus $\hat{\mu}$ lies on the boundary of $C(\Gamma)$. \square

3.4 2×2 Games

For a game Γ , the set $C(\Gamma)$, in general, need not be convex (Example 2 in [74]). In this section we will see that, in the special case of a 2×2 game with players having a fixed reference point independent of the underlying probability distribution, $C(\Gamma)$ is a convex polytope.

Consider a 2 player game Γ with $N = \{1, 2\}$ and $A_1 = A_2 = \{0, 1\}$. With player 1 as the row player and player 2 as the column player and $\{c_{ij}, d_{ij}\}_{i,j \in \{0,1\}}$ representing payoffs for player 1 and 2 respectively, let the payoff matrix be as shown in Figure 3.1. Here, the real numbers c_{ij} and d_{ij} should be thought of as outcomes in the terminology of cumulative prospect theory, but we will call them payoffs in this section. Let $\mu = \{\mu_{00}, \mu_{01}, \mu_{10}, \mu_{11}\} \in \Delta^3$ be a joint probability distribution assigning probabilities to action profiles as represented by the matrix in Figure 3.1. Let r_1 and r_2 be the fixed reference points (independent of the joint probability distribution μ) for players 1 and 2 respectively.

	0	1
0	c_{00}, d_{00}	c_{01}, d_{01}
1	c_{10}, d_{10}	c_{11}, d_{11}

	0	1
0	μ_{00}	μ_{01}
1	μ_{10}	μ_{11}

Figure 3.1: Payoff matrix (left) and joint probability matrix (right) of a 2×2 game

Proposition 3.4.1. *For the above 2×2 game, the set C_{CPT} is a convex polytope.*

Proof. The condition for $\mu \in C_{CPT}$ corresponding to the row player deviating from strategy 0 to strategy 1 in (3.2.3) is:

$$\text{if } \mu_{00} + \mu_{01} > 0 \text{ then } \mathcal{R}_1^{r_1}(p^1, z, y) \geq 0, \quad (3.4.1)$$

where $p^1 = (p_0^1, p_1^1)$, $p_0^1 = \frac{\mu_{00}}{\mu_{00} + \mu_{01}}$, $p_1^1 = \frac{\mu_{01}}{\mu_{00} + \mu_{01}}$, $z = (c_{00}, c_{01})$, $y = (c_{10}, c_{11})$. Let C_1 denote the set of all $\mu \in \Delta^3$ satisfying condition (3.4.1). We have:

- (i) if $c_{00} \geq c_{10}$ and $c_{01} \geq c_{11}$, then $C_1 = \Delta^3$;
- (ii) if $c_{00} < c_{10}$ and $c_{01} = c_{11}$ (resp. $c_{00} = c_{10}$ and $c_{01} < c_{11}$), then $C_1 = \{\mu \in \Delta^3 | \mu_{00} = 0\}$ (resp. $C_1 = \{\mu \in \Delta^3 | \mu_{01} = 0\}$);
- (iii) If $c_{00} < c_{10}$ and $c_{01} < c_{11}$, then $C_1 = \{\mu \in \Delta^3 | \mu_{00} = 0, \mu_{01} = 0\}$;

- (iv) if $c_{00} < c_{10}$ and $c_{01} > c_{11}$ (resp. $c_{00} > c_{10}$ and $c_{01} < c_{11}$), then from lemma 3.3.3, $\mathcal{R}_1^{r_1}(p^1, z, y)$ is strictly monotonic as a function of $p_0^1 (= 1 - p_1^1)$ on the interval $(0, 1)$,

$$\begin{aligned} \mathcal{R}_1^{r_1}((0, 1), z, y) &> 0 > \mathcal{R}_1^{r_1}((1, 0), z, y) \\ (\text{resp. } \mathcal{R}_1^{r_1}((0, 1), z, y) &< 0 < \mathcal{R}_1^{r_1}((1, 0), z, y)), \end{aligned}$$

and hence the inequality in condition (3.4.1) holds iff $p_0^1 \leq q_0$ (resp. $p_1^1 \leq q_1$) for a certain $q_0 \in (0, 1)$ (resp. $q_1 \in (0, 1)$) depending on the payoffs c_{00}, c_{01}, c_{10} and c_{11} , the value function $v_1^r(\cdot)$, and the probability weight functions $w_1^\pm(\cdot)$. Thus $C_1 = \{\mu \in \Delta^3 | \alpha_0 \mu_{00} \leq \mu_{01}\}$ with $\alpha_0 = \frac{1-q_0}{q_0}$ (resp. $C_1 = \{\mu \in \Delta^3 | \alpha_1 \mu_{00} \geq \mu_{01}\}$ with $\alpha_1 = \frac{q_1}{1-q_1}$).

In each case, C_1 is a convex polytope. Similarly, the other three conditions in (3.2.3), corresponding to the row player deviating from action 1 to action 0, the column player deviating from action 0 to action 1, and the column player deviating from action 1 to action 0, give rise to convex polytopes C_2, C_3 and C_4 respectively. The set $C(\Gamma)$ is the (non-empty) intersection of these convex polytopes and hence is itself a convex polytope. \square

Remark 3.4.2. From the assumption that the value functions and the probability weighting functions are continuous, we get that the CPT value function $V^r(p, z)$ is continuous in r, p and z [74], and hence $\mathcal{R}^{r_1}(p^1, z, y)$ in the proof above is continuous in r_1, p^1, z and y . In case (iv) above, since $R^{r_1}(p^1, z, y)$ is a strictly monotonic continuous function, the probability threshold q_0 (resp. q_1 depending on the relation between the payoffs) is uniquely determined by the payoff vectors x and y (keeping the reference point, value function and probability weighting functions fixed) as the one satisfying $\mathcal{R}^{r_1}(q_0, z, y) = 0$ (resp. $\mathcal{R}^{r_1}(q_1, z, y) = 0$). Let $(z^t = (c_{00}^t, c_{01}^t), t \geq 1)$ and $(y^t = (c_{10}^t, c_{11}^t), t \geq 1)$ be a sequence of payoff vectors such that $c_{00}^t < c_{10}^t$ and $c_{01}^t > c_{11}^t$ for all $t \geq 1$ and $z^t \rightarrow z^* = (c_{00}^*, c_{01}^*)$ and $y^t \rightarrow y^* = (c_{10}^*, c_{11}^*)$. Let q_0^t be the corresponding probability threshold for payoff vectors z^t and y^t . Unless $c_{00}^* = c_{10}^*$ and $c_{01}^* = c_{11}^*$, there is a unique q_0^* such that $\mathcal{R}^{r_1}(q_0^*, z^*, y^*) = 0$. If $q_0^t \not\rightarrow q_0^*$, then the sequence $(q_0^t, t \geq 1)$ has a limit point $\tilde{q}_0 \neq q_0^*$ and from the continuity of the function $\mathcal{R}^{r_1}(p, z, y)$ we get that $\mathcal{R}^{r_1}(\tilde{q}_0, z^*, y^*) = 0$ contradicting the uniqueness of q_0^* . Hence, except for the case when $c_{00}^* = c_{10}^*$ and $c_{01}^* = c_{11}^*$, we have $q_0^t \rightarrow q_0^*$ and hence $\alpha_0^t = \frac{1-q_0^t}{q_0^t} \rightarrow \alpha^* \in \mathbb{R}^+ \cup \{0, \infty\}$. We further note that the limit α^* depends only on the payoff vectors z^*, y^* and not on the limiting sequence z^t, y^t . This fact will be useful in analyzing 2×2 games with weakly dominated strategies as defined below. The case when $c_{00}^t > c_{10}^t$ and $c_{01}^t < c_{11}^t$ for all t is similar.

Definition 3.4.3. For an n player game $\Gamma = ([n], (A_i)_{i \in [n]}, (x_i)_{i \in [n]})$, let $a_i, \tilde{a}_i \in A_i$ be two strategies corresponding to player i .

- Strategies a_i and \tilde{a}_i are said to be equivalent if player i is indifferent in choosing between a_i and \tilde{a}_i no matter what the other players do.
- Strategy a_i is said to be weakly dominated by strategy \tilde{a}_i if there exists at least one strategy profile of the opponents for which choosing \tilde{a}_i is better than choosing a_i , and for all strategy profiles of the opponents choosing \tilde{a}_i is at least as good as choosing a_i .

- Strategy a_i is said to be strictly dominated by strategy \tilde{a}_i if, for every strategy profile of the opponents, choosing \tilde{a}_i is better than choosing a_i .

Note that a strictly dominated strategy is also a weakly dominated strategy.

As observed in Section 3.3, two outcome profiles z and y are equivalent under all probability distributions p iff $z = y$. Thus, as under EUT, for players with CPT preferences we have the following:

- Strategy a_i is equivalent to strategy \tilde{a}_i iff

$$x_i(a_i, a_{-i}) = x_i(\tilde{a}_i, a_{-i}) \quad \forall a_{-i} \in S_{-i}.$$

- Strategy a_i is weakly dominated by strategy \tilde{a}_i iff

$$x_i(a_i, a_{-i}) \leq x_i(\tilde{a}_i, a_{-i}) \quad \forall a_{-i} \in A_{-i},$$

where strict inequality holds for at least one $a_{-i} \in A_{-i}$.

- Strategy a_i is strictly dominated by strategy \tilde{a}_i iff

$$x_i(a_i, a_{-i}) < x_i(\tilde{a}_i, a_{-i}) \quad \forall a_{-i} \in S_{-i}.$$

We now look at the convex polytope $C(\Gamma)$ for a 2×2 game in more detail.

2×2 games with at least one equivalent pair of strategies

Consider a 2×2 game with at least one equivalent pair of strategies. Suppose player 1 has equivalent strategies. This corresponds to case (i) above with both equalities (i.e. $c_{00} = c_{10}$ and $c_{01} = c_{10}$). Thus player 1 is indifferent between his strategies. For player 2, if the two strategies are equivalent, then the game is trivial and $C(\Gamma) = \Delta^3$. If one of the strategies for player 2 is weakly dominated, say strategy 0 is weakly dominated by strategy 1, then either $d_{00} = d_{01}, d_{10} < d_{11}$ or $d_{00} < d_{01}, d_{10} = d_{11}$. If $d_{00} = d_{01}, d_{10} < d_{11}$, then the set $C(\Gamma) = \{\mu \in \Delta^3 | \mu_{10} = 0\}$ is a triangle with vertices $F = (1, 0, 0, 0)$, $G = (0, 1, 0, 0)$ and $H = (0, 0, 0, 1)$. It intersects the set $\Delta^*(A)$ at the lines with endpoints $\{F, G\}$ and $\{G, H\}$. The other three cases are similar. If neither of the two strategies for player 2 dominates the other, then $C(\Gamma)$ is characterized by the inequalities

$$\{\beta\mu_{00} \geq \mu_{10}, \beta\mu_{01} \leq \mu_{11}\} \text{ or } \{\beta\mu_{00} \leq \mu_{10}, \beta\mu_{01} \geq \mu_{11}\}.$$

where the former pair holds if $d_{00} > d_{01}, d_{10} < d_{11}$ and the latter holds if $d_{00} < d_{01}, d_{10} > d_{11}$. Suppose the first pair of inequalities hold (the other case can be handled similarly). Then one can check that the set $C(\Gamma)$ is a tetrahedron with vertices $E = (1, 0, 0, 0)$, $F = (\frac{1}{1+\beta}, 0, \frac{\beta}{1+\beta}, 0)$, $G = (0, 0, 0, 1)$ and $H = (0, \frac{1}{1+\beta}, 0, \frac{\beta}{1+\beta})$. It intersects $\Delta^*(A)$ at the lines with endpoints $\{E, F\}$, $\{G, H\}$ and $\{F, H\}$.

2×2 games with at least one strictly dominated strategy

Consider now a 2×2 game with at least one strictly dominated strategy. This corresponds to case (i) above with both inequalities strict (i.e. $c_{00} > c_{10}$ and $c_{01} > c_{11}$) or case (iii). Strictly dominated strategies cannot be used with positive probability in any correlated equilibrium of that game. It is easy then to compute the set $C(\Gamma)$ for such a game by eliminating the strictly dominated strategies. Suppose strategy 1 is strictly dominated by strategy 0. Thus, $c_{00} > c_{10}$ and $c_{01} > c_{11}$. If $d_{00} > d_{01}$ then $C(\Gamma) = \{\mu \in \Delta^3 | \mu_{01} = \mu_{10} = \mu_{11} = 0\}$ is a point. If $d_{00} < d_{01}$ then $C(\Gamma) = \{\mu \in \Delta^3 | \mu_{00} = \mu_{10} = \mu_{11} = 0\}$ is a point. If $d_{00} = d_{01}$ then $C(\Gamma) = \{\mu \in \Delta^3 | \mu_{10} = \mu_{11} = 0\}$ is a line segment. In each case $C(\Gamma)$ is contained in $\Delta^*(A)$. The case when strategy 0 is strictly dominated by strategy 1 is similar.

2×2 games with no equivalent or weakly dominated strategies

We now discuss 2×2 games with no equivalent or weakly dominated strategies. Let G^0 denote the set of all such games. For any game $\Gamma \in G^0$, the relation amongst the payoffs for all the four conditions corresponding to C_1, C_2, C_3 and C_4 are as in case (iv) above. Further, the conditions corresponding to the row player deviating from strategy 0 to strategy 1, and vice versa are

$$\text{if } \mu_{00} + \mu_{01} > 0 \text{ then } V_1^{r1}(p^1, z) \geq V_1^{r1}(p^1, y); \quad (3.4.2)$$

and

$$\text{if } \mu_{10} + \mu_{11} > 0 \text{ then } V_1^{r1}(p^2, z) \leq V_1^{r1}(p^2, y); \quad (3.4.3)$$

respectively, where p^1 is as in Proposition 3.4.1 and $p^2 = (p_0^2, p_1^2)$, $p_0^2 = \frac{\mu_{10}}{\mu_{10} + \mu_{11}}$ and $p_1^2 = \frac{\mu_{11}}{\mu_{10} + \mu_{11}}$. Now there exists a $q_0 \in (0, 1)$ (or a $q_1 \in (0, 1)$) such that inequality (3.4.2) holds for all $p_0^1 \leq q_0$ (resp. $p_1^1 \leq q_1$) and inequality (3.4.3) holds for all $p_0^2 \geq q_0$ (resp. $p_1^2 \geq q_1$). Thus if $C_1 = \{\mu \in \Delta^3 | \alpha_0 \mu_{00} \leq \mu_{01}\}$ (resp. $C_1 = \{\mu \in \Delta^3 | \alpha_1 \mu_{00} \geq \mu_{01}\}$), then $C_2 = \{\mu \in \Delta^3 | \alpha_0 \mu_{10} \geq \mu_{11}\}$ (resp. $C_2 = \{\mu \in \Delta^3 | \alpha_1 \mu_{10} \leq \mu_{11}\}$). Similarly for player 2. Thus, depending on the relation amongst the payoffs, the conditions (3.2.3) take one of the following forms:

(I) if $c_{00} > c_{10}, c_{01} < c_{11}, d_{00} > d_{01}, d_{10} < d_{11}$ then

$$\alpha \mu_{00} \geq \mu_{01}, \alpha \mu_{10} \leq \mu_{11}, \beta \mu_{00} \geq \mu_{10}, \beta \mu_{01} \leq \mu_{11};$$

(II) if $c_{00} < c_{10}, c_{01} > c_{11}, d_{00} > d_{01}, d_{10} < d_{11}$ then

$$\alpha \mu_{00} \leq \mu_{01}, \alpha \mu_{10} \geq \mu_{11}, \beta \mu_{00} \geq \mu_{10}, \beta \mu_{01} \leq \mu_{11};$$

(III) if $c_{00} > c_{10}, c_{01} < c_{11}, d_{00} < d_{01}, d_{10} > d_{11}$ then

$$\alpha \mu_{00} \geq \mu_{01}, \alpha \mu_{10} \leq \mu_{11}, \beta \mu_{00} \leq \mu_{10}, \beta \mu_{01} \geq \mu_{11};$$

	0	1
0	α, β	$0, 0$
1	$0, 0$	$1, 1$

$\gamma_I(\alpha, \beta)$

	0	1
0	$\alpha, -\beta$	$0, 0$
1	$0, 0$	$1, -1$

$\gamma_{III}(\alpha, \beta)$

	0	1
0	$-\alpha, \beta$	$0, 0$
1	$0, 0$	$-1, 1$

$\gamma_{II}(\alpha, \beta)$

	0	1
0	$-\alpha, -\beta$	$0, 0$
1	$0, 0$	$-1, -1$

$\gamma_{IV}(\alpha, \beta)$

Figure 3.2: Canonical 2×2 games

(IV) if $c_{00} < c_{10}, c_{01} > c_{11}, d_{00} < d_{01}, d_{10} > d_{11}$ then

$$\alpha\mu_{00} \leq \mu_{01}, \alpha\mu_{10} \geq \mu_{11}, \beta\mu_{00} \leq \mu_{10}, \beta\mu_{01} \geq \mu_{11};$$

for some $\alpha, \beta > 0$. Thus every 2×2 game with no equivalent or weakly dominated strategies can be classified into one of the above four types depending on the relations amongst its payoffs.

We consider the canonical 2×2 games $\gamma_l(\alpha, \beta)$ for $l \in \{I, II, III, IV\}$ with $\alpha, \beta > 0$ as shown in Figure 3.2. One can check that the set C_{EUT} for each of these games is given by the corresponding inequalities above.

As in the paper [23], based on the type of inequalities satisfied, we classify all 2×2 games, with no equivalent or weakly dominated strategies, into three types:

- coordination games if the inequalities take form (I),
- anti-coordination games if the inequalities take form (IV) and,
- competitive games if the inequalities take either form (II) or form (III).

Since the inequalities above completely characterize the set $C(\Gamma)$, it is enough to find the set $C_{EUT}(\Gamma)$ for each of the canonical games. For case (II), we have

$$\alpha\mu_{00} \leq \mu_{01} \leq \frac{\mu_{11}}{\beta} \text{ and } \mu_{11} \leq \alpha\mu_{10} \leq \beta\alpha\mu_{00}.$$

μ	μ_{00}	μ_{01}	μ_{10}	μ_{11}
$\mu_A^*(\alpha, \beta)$	1	0	0	0
$\mu_B^*(\alpha, \beta)$	0	0	0	1
$\mu_C^*(\alpha, \beta)$	$\frac{1}{(1+\alpha)(1+\beta)}$	$\frac{\alpha}{(1+\alpha)(1+\beta)}$	$\frac{\beta}{(1+\alpha)(1+\beta)}$	$\frac{\alpha\beta}{(1+\alpha)(1+\beta)}$
$\mu_D^*(\alpha, \beta)$	$\frac{1}{1+\beta+\alpha\beta}$	0	$\frac{\beta}{1+\beta+\alpha\beta}$	$\frac{\alpha\beta}{1+\beta+\alpha\beta}$
$\mu_E^*(\alpha, \beta)$	$\frac{1}{1+\alpha+\alpha\beta}$	$\frac{\alpha}{1+\alpha+\alpha\beta}$	0	$\frac{\alpha\beta}{1+\alpha+\alpha\beta}$

Figure 3.3: Vertices of the convex polytope C_{EUT} for $\gamma_I(\alpha, \beta)$

Thus all inequalities must be satisfied with equality and we get

$$\begin{aligned}\mu_{00} &= \frac{1}{(1+\alpha)(1+\beta)}, \mu_{01} = \frac{\alpha}{(1+\alpha)(1+\beta)}, \\ \mu_{10} &= \frac{\beta}{(1+\alpha)(1+\beta)}, \mu_{11} = \frac{\alpha\beta}{(1+\alpha)(1+\beta)}.\end{aligned}\tag{3.4.4}$$

Case (III) is similar. Thus, for competitive games, the set $C_{EUT}(\Gamma)$ is reduced to a single point, which is also the unique mixed Nash equilibrium. For coordination games, the set $C_{EUT}(\Gamma)$ is a convex polytope with five vertices as given in Figure 3.3. It intersects the set $\Delta^*(A)$ at the three vertices $\mu_A^*(\alpha, \beta)$, $\mu_B^*(\alpha, \beta)$ and $\mu_C^*(\alpha, \beta)$ of which the first two are pure Nash equilibria. From the set of inequalities corresponding to cases (I) and (IV) we can see that the joint distribution $\mu = (\mu_{00}, \mu_{01}, \mu_{10}, \mu_{11})$ belongs to $C_{EUT}(\Gamma)$ of $\gamma_I(\alpha, \beta)$ iff $\tau(\mu) := (\mu_{10}, \mu_{11}, \mu_{00}, \mu_{01})$ belongs to $C_{EUT}(\Gamma)$ of $\gamma_{IV}(\alpha, 1/\beta)$. Thus, for anti-coordination games, the set $C_{EUT}(\Gamma)$ is again a convex polytope with five vertices and it intersects $\Delta^*(A)$ at three of its vertices with two of them pure Nash equilibria. The vertices can be found from Figure 3.3, using the transformation τ , and replacing β by $1/\beta$. Since $C(\Gamma)$ is determined by the same set of inequalities in α and β , all the results carry over to 2×2 games with CPT preferences. In particular, the set $C(\Gamma)$ is determined by α and β . Since the set of CPT Nash equilibria (pure and mixed) is given by the intersection of $\Delta^*(A)$ and $C(\Gamma)$, we have a unique mixed CPT Nash equilibrium and no pure CPT Nash equilibria for competitive games, and one mixed and two pure CPT Nash equilibria for coordination and anti-coordination games.

2×2 games with at least one weakly dominated strategy but no equivalent or strictly dominated strategy

Let G^1 denote the set of all 2×2 games with at least one weakly dominated strategy but no equivalent or strictly dominated strategy. If player 1 has a weakly dominated strategy then this corresponds to case (i) above with one equality and one strict inequality (i.e. $c_{00} > c_{10}$ and $c_{01} = c_{11}$, or $c_{00} = c_{10}$ and $c_{01} > c_{11}$), or case (ii). The set of all 2×2 games, each game characterized by its payoff matrix, forms an 8-dimensional Euclidean vector space.

	$\beta = 0$	$0 < \beta < \infty$	$\beta = \infty$
$\alpha = 0$	$\mu_{01} = 0$ $\mu_{10} = 0$	$\mu_{01} = 0$ $\beta\mu_{00} \geq \mu_{10}$	$\mu_{01} = 0$
$0 < \alpha < \infty$	$\mu_{10} = 0$ $\alpha\mu_{00} \geq \mu_{01}$	–	$\mu_{01} = 0$ $\alpha\mu_{10} \leq \mu_{11}$
$\alpha = \infty$	$\mu_{10} = 0$	$\mu_{10} = 0$ $\beta\mu_{01} \leq \mu_{11}$	$\mu_{01} = 0$ $\mu_{10} = 0$

Figure 3.4: The set C_{CPT} for games of type (I) with weakly dominated strategies

For every game $\Gamma \in G^1$, the intersection of any ϵ -neighborhood of Γ for sufficiently small $\epsilon > 0$ with the set G^0 is non-empty and contains games of a unique type and hence Γ can be seen as a limit of games in G^0 of a unique type $l \in \{I, II, III, IV\}$.¹¹ Let $\alpha^l(\tilde{\Gamma}), \beta^l(\tilde{\Gamma})$, for $\tilde{\Gamma} \in G^0$, be such that $\gamma_l(\alpha^l(\tilde{\Gamma}), \beta^l(\tilde{\Gamma}))$ are the corresponding canonical games to $\tilde{\Gamma}$. From Remark 3.4.2, for any sequence $(\Gamma_t, t \geq 1)$ of games $\Gamma_t \in G^0$ of this unique type l , such that $\Gamma_t \rightarrow \Gamma$, the sequences $\alpha^l(\Gamma_t) \rightarrow \alpha$ and $\beta^l(\Gamma_t) \rightarrow \beta$, where α, β depend on the game Γ and not on the sequence $(\Gamma_t, t \geq 1)$.

For example, a game $\Gamma \in G^1$ with payoffs satisfying $c_{00} < c_{10}, c_{01} = c_{11}, d_{00} > d_{01}, d_{10} < d_{11}$ is the limit, as $\epsilon \downarrow 0$, of games $\Gamma_\epsilon \in G^0$ with payoffs same as that of Γ except c_{01} replaced by $c_{01} + \epsilon$. Each of the games Γ_ϵ is of type (II) for sufficiently small $\epsilon > 0$. Further, if $\gamma_{II}(\alpha_\epsilon, \beta_\epsilon)$ are the canonical games corresponding to Γ_ϵ then $\alpha_\epsilon \rightarrow \alpha = \infty$ and $\beta_\epsilon = \beta$ for some fixed β .

Suppose $\gamma_l(\alpha, \beta)$ is the corresponding canonical game for Γ . Then the payoffs of player 1 are strategically equivalent to

$$M = \begin{bmatrix} \alpha & 0 \\ 0 & 1 \end{bmatrix}.$$

Strategically equivalent here should be interpreted as meaning that the best response of player 1 to a mixed strategy $p^1 = (p_0^1, p_1^1)$ in a game, with her payoffs given by matrix M under EUT, is same as that in the game Γ under CPT. When $\alpha = \infty$, it means that action 1 is player 1's best response only to the mixed strategy $p^1 = (0, 1)$ and action 0 is her best response for all mixed strategies p^1 , which is true when action 0 weakly dominates action 1 but not strictly.

Using this observation we classify every game $\Gamma \in G^1$ into four types $l \in \{I, II, III, IV\}$ and each of these 4 types into eight further subtypes depending on the limit of α_ϵ and β_ϵ going either to $0, \infty$ or some real number in $(0, \infty)$ with at least one of them tending to 0 or ∞ . The set $C(\Gamma)$ for games of the type (I) and (II) are given in Figures 3.4 and 3.5 respectively. The set $C(\Gamma)$ for types (IV) and (III) can be found using the transformation τ and replacing β by $1/\beta$ (with the convention $1/0 = \infty$ and $1/\infty = 0$) from Figures 3.4 and 3.5 respectively.

	$\beta = 0$	$0 < \beta < \infty$	$\beta = \infty$
$\alpha = 0$	$\mu_{11} = 0$ $\mu_{10} = 0$	$\mu_{11} = \mu_{01} = 0$ $\beta\mu_{00} \geq \mu_{10}$	$\mu_{11} = 0$ $\mu_{01} = 0$
$0 < \alpha < \infty$	$\mu_{11} = \mu_{10} = 0$ $\alpha\mu_{00} \leq \mu_{01}$	—	$\mu_{01} = \mu_{00} = 0$ $\alpha\mu_{10} \geq \mu_{11}$
$\alpha = \infty$	$\mu_{10} = 0$ $\mu_{00} = 0$	$\mu_{00} = \mu_{10} = 0$ $\beta\mu_{01} \leq \mu_{11}$	$\mu_{00} = 0$ $\mu_{01} = 0$

Figure 3.5: The set C_{CPT} for games of type (II) with weakly dominated strategies

We now describe the typical geometry displayed by the CPT equilibrium notions in each of the above cases. The geometry of C_{CPT} in case (I) is as follows:

- If $\alpha = \beta = 0$ or $\alpha = \beta = \infty$, then $C(\Gamma)$ is a line with endpoints $F = (1, 0, 0, 0)$ and $G = (0, 0, 0, 1)$. It intersects the set I at the two endpoints F and G .
- If $\alpha = 0, \beta = \infty$, then $C(\Gamma)$ is a triangle with vertices $F = (1, 0, 0, 0)$, $G = (0, 0, 1, 0)$, and $H = (0, 0, 0, 1)$. It intersects $\Delta^*(A)$ at the lines with endpoints $\{F, G\}$ and $\{G, H\}$. Similarly, if $\alpha = \infty, \beta = 0$, then C_{CPT} is a triangle and it intersects $\Delta^*(A)$ at two lines.
- If $\alpha = 0, 0 < \beta < \infty$, then $C(\Gamma)$ is a triangle with vertices $F = (0, 0, 0, 1)$, $G = (1, 0, 0, 0)$, and $H = (\frac{1}{1+\beta}, 0, \frac{\beta}{1+\beta}, 0)$. It intersects the set $\Delta^*(A)$ at the point $(0, 0, 0, 1)$ and the line joining the points $(1, 0, 0, 0)$ and $(\frac{1}{1+\beta}, 0, \frac{\beta}{1+\beta}, 0)$. The remaining three cases can be analyzed similarly.

The geometry of $C(\Gamma)$ in case (II) is as follows:

- If $\alpha = \beta = 0$, then the set $C(\Gamma)$ is a line joining the points $(1, 0, 0, 0)$ and $(0, 1, 0, 0)$ contained in the set $\Delta^*(A)$. Similarly, for the cases when $\alpha = 0$ or ∞ and $\beta = 0$ or ∞ , the set $C(\Gamma)$ is a line segment contained in the set $\Delta^*(A)$.
- If $\alpha = 0, 0 < \beta < \infty$, then $C(\Gamma)$ is a line joining the points $(1, 0, 0, 0)$ and $(\frac{1}{1+\beta}, 0, \frac{\beta}{1+\beta}, 0)$ and is contained in the set I . The remaining three cases can be analyzed similarly.

The geometry of $C(\Gamma)$ in cases (IV) and (III) can be obtained from cases (I) and (II) respectively, using the transformation τ , and replacing β by $1/\beta$.

3.5 On the connectedness of CPT correlated equilibrium

In the previous section, we saw that for 2×2 games the set $C(\Gamma)$ is a convex polytope. However, in general, the set $C(\Gamma)$ can have a more complicated geometry. We will now see that the set C_{CPT} can, in fact, be disconnected.

In this section, we restrict our attention to games with each player i having reference point $r_i = 0$, and all the outcomes $x_i(\cdot)$ non-negative. Thus all our outcome profiles are “one-sided” with zero reference point, and we will denote $w_i^+(\cdot)$, $v_i^0(\cdot)$, $V_i^r(\cdot)$ simply by $w_i(\cdot)$, $v_i(\cdot)$, $V_i(\cdot)$ respectively.

The geometry of the set C_{CPT} is determined by the set of inequalities (3.2.3). Let us consider the inequality corresponding to player i deviating from strategy a_i to \tilde{a}_i . For ease of notation, fix a one to one correspondence between the numbers $\{1, \dots, t\}$ and the joint strategies $\{a_{-i} \in A_{-i}\}$ (here $t = |A_{-i}|$). Let

$$z = (z_1, \dots, z_t) := (v_i(x_i(a_i, a_{-i})))_{a_{-i} \in A_{-i}},$$

and

$$y = (y_1, \dots, y_t) := (v_i(x_i(\tilde{a}_i, a_{-i})))_{a_{-i} \in A_{-i}}.$$

Let $p = (p_1, \dots, p_t) \in \Delta^{t-1}$ be a joint probability distribution on S_{-i} . Let $(\alpha_1, \dots, \alpha_t)$ and $(\beta_1, \dots, \beta_t)$ be permutations of $(1, \dots, t)$ such that

$$z_{\alpha_1} \geq z_{\alpha_2} \geq \dots \geq z_{\alpha_t} \text{ and } y_{\beta_1} \geq y_{\beta_2} \geq \dots \geq y_{\beta_t},$$

respectively.

Consider the inequality

$$\tilde{V}_i(p, z) \geq \tilde{V}_i(p, y), \quad (3.5.1)$$

where

$$\begin{aligned} \tilde{V}_i(p, z) &= z_{\alpha_t} + w_i(p_{\alpha_1} + \dots + p_{\alpha_{t-1}})[z_{\alpha_{t-1}} - z_{\alpha_t}] \\ &\quad + w_i(p_{\alpha_1} + \dots + p_{\alpha_{t-2}})[z_{\alpha_{t-2}} - z_{\alpha_{t-1}}] + \dots + w_i(p_{\alpha_1})[z_{\alpha_1} - z_{\alpha_2}], \end{aligned} \quad (3.5.2)$$

and

$$\begin{aligned} \tilde{V}_i(p, y) &= y_{\beta_t} + w_i(p_{\beta_1} + \dots + p_{\beta_{t-1}})[y_{\beta_{t-1}} - y_{\beta_t}] \\ &\quad + w_i(p_{\beta_1} + \dots + p_{\beta_{t-2}})[y_{\beta_{t-2}} - y_{\beta_{t-1}}] + \dots + w_i(p_{\beta_1})[y_{\beta_1} - y_{\beta_2}]. \end{aligned} \quad (3.5.3)$$

To contrast with the notation used in earlier sections, note that $\tilde{V}_i(p, z) = V_i(p, \tilde{z})$ and $\tilde{V}_i(p, y) = V_i(p, \tilde{y})$, where $\tilde{x} := (x_i(a_i, a_{-i}))_{a_{-i} \in A_{-i}}$ and $\tilde{y} := (x_i(\tilde{a}_i, a_{-i}))_{a_{-i} \in A_{-i}}$. Let $C(\Gamma, i, a_i, \tilde{a}_i)$ denote the set of all probability vectors $p \in \Delta^{t-1}$ that satisfy the inequality (3.5.1). We can similarly define $C(\Gamma, i, s_i, d_i)$ for all $i \in [n]$, $a_i, \tilde{a}_i \in A_i$. It is clear from the definition of CPT correlated equilibrium that for a joint probability distribution $\mu \in C(\Gamma)$, provided $\mu_i(s_i) > 0$, the probability vector $p = \mu_{-i}(\cdot | a_i) \in \Delta^{t-1}$ should belong to $C(\Gamma, i, a_i, \tilde{a}_i)$ for all $\tilde{a}_i \in A_i$. Let

$$C(\Gamma, i, a_i) := \bigcap_{\tilde{a}_i \in A_i} C(\Gamma, i, a_i, \tilde{a}_i).$$

Now, for all i , define a subset $C(\Gamma, i) \subset \Delta^{|A|-1}$, as follows:

$$C(\Gamma, i) := \{\mu \in \Delta^{|A|-1} \mid \mu_{-i}(\cdot | a_i) \in C(\Gamma, i, a_i), \forall a_i \in A_i \text{ such that } \mu_i(a_i) > 0\}.$$

Note that since $C(\Gamma)$ is nonempty, the set $C(\Gamma, i)$ is nonempty for each i . The set $C(\Gamma, i)$ can be constructed from the sets $\{C(\Gamma, i, a_i), a_i \in A_i\}$ as follows: let $p^{a_i} \in C(\Gamma, i, a_i)$ for all $a_i \in A_i$ such that $C(\Gamma, i, a_i) \neq \emptyset$, let $q_i \in \Delta^{|A_i|-1}$ be a probability distribution over A_i such that $q_i(a_i) = 0$ for all $a_i \in S_i$ such that $C(\Gamma, i, a_i) = \emptyset$, and define a joint probability distribution $\mu \in \Delta^{|A|-1}$ by $\mu(a_i, a_{-i}) = q_i(a_i)p^{a_i}(a_{-i})$ if $C(\Gamma, i, a_i) \neq \emptyset$ and $\mu(a_i, a_{-i}) = 0$ otherwise. Then $\mu \in C(\Gamma, i)$, and for every $\mu \in C(\Gamma, i)$, the corresponding $q_i = \mu_i$ for all $a_i \in A_i$ and $p^{a_i} = \mu_{-i}(\cdot|a_i)$ for all $a_i \in A_i$ with $C(\Gamma, i, a_i) \neq \emptyset$. Further, it is clear that

$$C(\Gamma) = \bigcap_{i \in [n]} C(\Gamma, i).$$

Thus the set $C(\Gamma)$ is uniquely determined by the collection of sets

$$\{C(\Gamma, i, a_i, \tilde{a}_i), i \in [n], a_i, \tilde{a}_i \in A_i\}.$$

Lemma 3.5.1. *In the above setting, the set $C(\Gamma, i, a_i, \tilde{a}_i)$ is connected.*

Proof. Suppose the permutations $(\alpha_1, \dots, \alpha_t)$ and $(\beta_1, \dots, \beta_t)$ can be chosen such that they are equal. Let

$$l_j := w_i\left(\sum_{k=1}^j p_{\alpha_k}\right) = w_i\left(\sum_{k=1}^j p_{\beta_k}\right), \text{ for } 1 \leq j \leq t. \quad (3.5.4)$$

For every vector $l = (l_1, \dots, l_t) \in \mathbb{R}^t$ such that $0 \leq l_1 \leq \dots \leq l_t = 1$, there corresponds a unique probability vector $p = (p_1, \dots, p_t)$ satisfying equations (3.5.4) and this mapping is continuous because $w_i(\cdot)$ is a continuous strictly increasing function. Thus we have a one-to-one correspondence between probability vectors (p_1, \dots, p_t) and the vectors (l_1, \dots, l_t) .

Inequality (3.5.1) can then be written as

$$l_t z_{\alpha_t} + \sum_{i=1}^{t-1} l_{t-i} [z_{\alpha_{t-i}} - z_{\alpha_{t-i+1}}] \geq l_t y_{\beta_t} + \sum_{i=1}^{t-1} l_{t-i} [y_{\beta_{t-i}} - y_{\beta_{t-i+1}}]. \quad (3.5.5)$$

Since this is linear in (l_1, \dots, l_t) , the set of all vectors (l_1, \dots, l_t) satisfying inequality (3.5.5) is a convex polytope. In particular, it is connected. Thus the set $C(\Gamma, i, a_i, \tilde{a}_i)$ is also connected.

Suppose now the permutations $(\alpha_1, \dots, \alpha_t)$ and $(\beta_1, \dots, \beta_t)$ cannot be chosen to be equal. Then there exists $1 \leq j_1, j_2 \leq t$ such that $z_{j_1} > z_{j_2}$ and $y_{j_1} \leq y_{j_2}$. If $p \in C(\Gamma, i, a_i, \tilde{a}_i)$ such that $p_{j_2} > 0$, then, by the stochastic dominance property, the following probability vector $q(\epsilon)$, for all $0 \leq \epsilon \leq 1$, also belongs to $C(\Gamma, i, a_i, \tilde{a}_i)$:

$$q_j(\epsilon) = \begin{cases} p_{j_1} + (1 - \epsilon)p_{j_2} & \text{if } j = j_1, \\ \epsilon p_{j_2} & \text{if } j = j_2, \\ p_j & \text{otherwise.} \end{cases}$$

	RED	YELLOW	GREEN
TOP	69, 10	61, 0	20, 10
CENTER	50, 0	60, 10	30, 0
BOTTOM	101, 0	41, 10	0, 0

Table 3.1: Payoff matrix for the game in Section 3.6

Thus, from every vector $p \in C(\Gamma, i, a_i, \tilde{a}_i)$, we have a path connecting it to a probability vector $p' \in C(\Gamma, i, a_i, \tilde{a}_i)$ such that $p'_{j_2} = 0$. To show that $C(\Gamma, i, a_i, \tilde{a}_i)$ is connected it is enough to show that the subset

$$C'(\Gamma, i, a_i, \tilde{a}_i) = \{p' \in C(\Gamma, i, a_i, \tilde{a}_i) | p'_{j_2} = 0\}.$$

is connected. From (3.5.2) and (3.5.3), we can see that the CPT values of the prospects (p, z) and (p, y) with probability vector restricted to $C'(\Gamma, i, a_i, \tilde{a}_i)$ do not depend on the outcomes z_{j_2} and y_{j_2} . If one can now choose permutations $(\alpha'_1, \dots, \alpha'_{t-1})$ and $(\beta'_2, \dots, \beta'_{t-1})$ of $\{1, \dots, t\} \setminus \{j_2\}$ such that

$$z_{\alpha'_1} \geq z_{\alpha'_2} \geq \dots \geq z_{\alpha'_{t-1}} \text{ and } y_{\beta'_1} \geq y_{\beta'_2} \geq \dots \geq y_{\beta'_{t-1}},$$

then, as before, one can argue that the set $C'(i, s_i, d_i)$ is connected. If not, we can continue to decrease the support of the probability vectors under consideration. This process terminates since our state space is finite. \square

Even though the sets $C(\Gamma, i, a_i, \tilde{a}_i)$ are connected, their intersection might be disconnected, as in the example given in the next section.

3.6 Example of a Game with Disconnected CPT Correlated Equilibrium

Consider a 2 player Γ game with each player having three pure strategies: TOP, CENTER, BOTTOM for player 1 (row player) and RED, YELLOW, GREEN for player 2 (column player), with the corresponding payoffs as shown in Table 3.1. For both the players, let $v_i(\cdot)$ be the identity function. For the probability weight function $w_i(\cdot)$ we employ the function suggested by Prelec [113], which, for $i = 1, 2$, is given by

$$w_i(p) = \exp\{-(-\ln p)^{\gamma_i}\},$$

for some $\gamma_i \in (0, 1]$. We take $\gamma_1 = 0.5$ and $\gamma_2 = 1$. We will now see that the set $C(\Gamma, 1, \text{TOP})$ is disconnected. Fix the correspondence $(R, Y, G) \leftrightarrow (\text{RED}, \text{YELLOW}, \text{GREEN})$. The set

$C(1, \text{TOP}, \text{BOTTOM})$ consists of all probability vectors $p = (p_R, p_Y, p_G) \in \Delta^2$ satisfying the following inequality:

$$\begin{aligned} 20 + w_1(p_R + p_Y)[61 - 20] + w_1(p_R)[69 - 61] \\ \geq 0 + w_1(p_R + p_Y)[41 - 0] + w_1(p_R)[101 - 41]. \end{aligned}$$

This holds iff $p_R \leq 0.40$ (all the decimal numbers henceforth are correct up to two decimal points). Thus, we have

$$C(\Gamma, 1, \text{TOP}, \text{BOTTOM}) = \{p \in \Delta^2 \mid p_R \leq 0.40\}.$$

The set $C(\Gamma, 1, \text{TOP}, \text{CENTER})$ consists of all probability vectors $p = (p_R, p_Y, p_G) \in \Delta^2$ satisfying the inequality

$$\begin{aligned} 20 + w_1(p_R + p_Y)[61 - 20] + w_1(p_R)[69 - 61] \\ \geq 30 + w_1(p_R + p_Y)[50 - 30] + w_1(p_Y)[60 - 50]. \end{aligned}$$

Rearranging, we get

$$21w_1(1 - p_G) - 10w_1(1 - p_R - p_G) \geq 10 - 8w_1(p_R).$$

For each $p_R \in [0, 0.4]$, we solve the above inequality for p_G . The set $C(\Gamma, 1, \text{TOP})$, as shown in Figure 3.1, is disconnected. One can check that

$$(0, \epsilon, 1 - \epsilon) \in C(\Gamma, 1, \text{CENTER}) \text{ and } (1 - \epsilon, \epsilon, 0) \in C(\Gamma, 1, \text{BOTTOM}),$$

for $\epsilon \in [0, 0.20]$. We cannot as yet conclude that the set $C(\Gamma, 1)$ is disconnected, because of the existence of joint probability distributions μ with marginal distribution $\mu_1(\text{TOP}) = 0$. We now show that $C(\Gamma, 2)$ cannot contain any distribution μ with $\mu_1(\text{TOP}) = 0$.

Fix the correspondence $(T, C, B) \leftrightarrow (\text{TOP}, \text{CENTER}, \text{BOTTOM})$. A similar analysis for player 2 shows that

$$\begin{aligned} C(\Gamma, 2, \text{RED}) &= \{p \in \Delta^2 \mid p_T \geq 0.5\}, \\ C(\Gamma, 2, \text{YELLOW}) &= \{p \in \Delta^2 \mid p_T \leq 0.5\}, \\ C(\Gamma, 2, \text{GREEN}) &= \{p \in \Delta^2 \mid p_T \geq 0.5\}. \end{aligned}$$

Suppose that there were $\mu \in C(\Gamma)$ with $\mu_1(\text{TOP}) = 0$. Then

$$\mu(\text{TOP}, \text{RED}) = \mu(\text{TOP}, \text{YELLOW}) = \mu(\text{TOP}, \text{GREEN}) = 0,$$

and from the structure of the sets $C(2, \text{RED})$ and $C(2, \text{GREEN})$ we get

$$\mu_2(\text{RED}) = \mu_2(\text{GREEN}) = 0.$$




Figure 3.1: Standard 2-simplex of probability vectors $p = (p_R, p_Y, p_G)$. The shaded region represents the set $C(\Gamma, 1, \text{TOP})$ and is disconnected.

Thus, the joint probability μ has support only on the strategy pairs (CENTER, YELLOW) and (BOTTOM, YELLOW). Thus, player 2 always plays strategy YELLOW and player 1 mixes between CENTER and BOTTOM. However, given player 2 plays strategy YELLOW, player 1's TOP strategy dominates strategies CENTER and BOTTOM. Hence such a joint probability distribution is not possible. Thus there does not exist any distribution $\mu \in C(\Gamma)$ with $\mu_1(\text{TOP}) = 0$.

There is a possibility that one of the components of $C(\Gamma, 1, \text{TOP})$ could disappear in the intersection $C(\Gamma, 1) \cap C(2)$. However, this does not happen because both the distributions $\bar{\mu}, \tilde{\mu}$ in Figure 3.2 belong to $C(\Gamma)$ with $\bar{\mu}_{-1}(\cdot | \text{TOP})$ and $\tilde{\mu}_{-1}(\cdot | \text{TOP})$ belonging to different components of $C(\Gamma, 1, \text{TOP})$. \square

3.7 Summary

Although the set of correlated equilibria under CPT has a more complicated geometry than a convex polytope, property (P), on the intersection of the Nash and correlated equilibrium sets, continues to hold. Property (P) is particularly relevant to the interactive learning problem in game theory [50, 51, 60]. This raises the question of analyzing the interactive learning problem under cumulative prospect theoretic preferences. We will get back to this in Chapter 5. In this process, we will see that the notion of CPT correlated equilibrium needs to be reconsidered.

	RED	YELLOW	GREEN
TOP	0.4	0.1	0.5
CENTER	0	0.05	0.5
BOTTOM	0.4	0.05	0

	RED	YELLOW	GREEN
TOP	0.4	0	0.6
CENTER	0	0	0.6
BOTTOM	0.4	0	0

Figure 3.2: Un-normalized distributions $\bar{\mu}$ and $\tilde{\mu}$.

In the next chapter, we will see another interesting phenomenon related to the notion of Nash equilibria when players have CPT preferences. We will rethink the reasoning behind the definition of Nash equilibrium and discuss its implications to the definition of CPT Nash equilibrium. This will give rise to two novel notions of equilibrium that we call black-box equilibrium.

Notes

⁹A working paper [23] analyzes the set of correlated equilibria for 2x2 games. However, some of the results in that paper are partially incorrect, and unfortunately, due to the untimely death of the author, they never got addressed. Our characterization is based on the ideas presented in that paper.

¹⁰Keskin defines CPT equilibrium assuming $w^+(\cdot) = w^-(\cdot)$. However, the definition can be easily extended to our general setting and the proof of existence goes through without difficulty.

¹¹Even though the game Γ can be obtained as a limit of a sequence of games in G^0 , one cannot obtain, in general, the set $C(\Gamma)$ as a limit of the sets of correlated equilibria of the games in G^0 . This is because the set of correlated equilibria is only upper-semicontinuous as a function of the game.

Chapter 4

Black-Box Equilibrium: Reconsidering CPT Nash Equilibrium

4.1 Introduction

The independence axiom says that if lottery L_1 is weakly preferred over lottery L_2 by an agent (i.e. the agent wants lottery L_1 at least as much as lottery L_2), and L is some other lottery, then, for $0 \leq \alpha \leq 1$, the combined lottery $\alpha L_1 + (1 - \alpha)L$ is weakly preferred over the combined lottery $\alpha L_2 + (1 - \alpha)L$ by that agent. A weakened form of the independence axiom, called betweenness, says that if lottery L_1 is weakly preferred over lottery L_2 (by an agent), then, for any $0 \leq \alpha \leq 1$, the mixed lottery $L = \alpha L_1 + (1 - \alpha)L_2$ must lie between the lotteries L_1 and L_2 in preference. Betweenness implies that if an agent is indifferent between L_1 and L_2 , then she is indifferent between any mixtures of them too. It is known that independence implies betweenness, but betweenness does not imply independence [34]. As a result, EUT preferences, which are known to satisfy the independence axiom, also satisfy betweenness. CPT preferences, on the other hand, do not satisfy betweenness in general (see Example 4.2.2). In fact, in Theorem 4.2.3, we show that CPT preferences satisfy betweenness if and only if they are EUT preferences (recall that EUT preferences are a special case of CPT preferences).

Suppose in a non-cooperative game that given her beliefs about the other players, a player is indifferent between two of her actions. Then according to EUT, she should be indifferent between any of the mixtures of these two actions. This facilitates the proof of the existence of a Nash equilibrium in mixed actions for such games. However, with CPT preferences, the player could either prefer some mixture of these two actions over the individual actions or vice versa.

As a result, it is important to make a distinction in CPT regarding whether the players can actively randomize over their actions or not. One way to enable active randomization is by assuming that each player has access to a randomizing device and the player can “commit” to the outcome of this randomization. The commitment assumption is necessary,

as is evident from the following scenario (the gambles presented below appear in [112]). Alice needs to choose between the following two actions:

1. Action 1 results in a lottery $L_1 = \{(0.34, \$20,000); (0.66, \$0)\}$, i.e. she receives \$20,000 with probability 0.34 and nothing with probability 0.66.
2. Action 2 results in a lottery $L_2 = \{(0.17, \$30,000); (0.83, \$0)\}$.

(See Example 4.3.7 for an instance of a 2-player game with Alice and Bob, where Alice has two actions that result in the above two lotteries.) Note that L_1 is a less risky gamble with a lower reward and L_2 is a more risky gamble with a higher reward. Now consider a compound lottery $L = 16/17L_1 + 1/17L_2$. Substituting for the lotteries L_1 and L_2 we get L in its reduced form to be

$$L = \{(0.01, \$30,000); (0.32, \$20,000); (0.67, \$0)\}.$$

In Example 4.2.1, we provide a CPT model for Alice's preferences that result in lottery L_1 being preferred over lottery L_2 , whereas lottery L is preferred over lotteries L_1 and L_2 . Roughly speaking, the underlying intuition is that Alice is risk-averse in general, and she prefers lottery L_1 over lottery L_2 . However, she over-weights the small 1% chance of getting \$30,000 in L and finds it lucrative enough to make her prefer lottery L over both the lotteries L_1 and L_2 . Let us say Alice has a biased coin that she can use to implement the randomized strategy. Now, if Alice tossed the coin, and the outcome was to play action 2, then in the absence of commitment, she will switch to action 1, since she prefers lottery L_1 over lottery L_2 . Commitment can be achieved, for example, by asking a trusted party to implement the randomized strategy for her or use a device that would carry out the randomization and implement the outcome without further consultation with Alice. Regardless of the implementation mechanism, we will call such randomized strategies *black-box strategies*. The above problem of commitment is closely related to the problem of using non-EUT models in dynamic decisions. For an interesting discussion on this topic, see Appendix C of [132] and the references therein.

Traditionally, mixed actions have been considered from two viewpoints, especially in the context of mixed action Nash equilibrium. According to the first viewpoint, these are conscious randomizations by the players – each player only knows her mixed action and not its pure realization. The notion of black-box strategies captures this interpretation of mixed actions. According to the other viewpoint, players do not randomize, and each player chooses some definite action, but the other players need not know which one, and the mixture represents their uncertainty, i.e. their conjecture about her choice. Aumann and Brandenburger [8] establish mixed action Nash equilibrium as an equilibrium in conjectures provided they satisfy certain epistemic conditions regarding the common knowledge amongst the players.

In the absence of the betweenness condition, these two viewpoints give rise to different notions of Nash equilibria. Throughout we assume that the player set and their corresponding action sets and payoff functions, as well as the rationality of each player, are common

knowledge. A player is said to be rational if, given her beliefs and her preferences, she does not play any suboptimal strategy. Suppose each player plays a fixed action, and these fixed actions are common knowledge, then we get back the notion of *pure Nash equilibrium* (see Definition 4.3.2). If each player plays a fixed action, but the other players have mixed conjectures over her action, and these conjectures are common knowledge, then this gives us *mixed action Nash equilibrium* (see Definition 4.3.4). This coincides with the notion of Nash equilibrium that we saw in Chapter 3. Now suppose each player can randomize over her actions and hence implement a black-box strategy. If each player plays a fixed black-box strategy and these black-box strategies are common knowledge, then this gives rise to a new notion of equilibrium. We call it *black-box strategy Nash equilibrium* (see Definition 4.3.8). If each player plays a fixed black-box strategy and the other players have mixed conjectures over her black-box strategy, and these conjectures are common knowledge, then we get the notion of *mixed black-box strategy Nash equilibrium* (see Definition 4.3.10). It should be noted that the notion of mixed black-box strategy Nash equilibrium is identical to the notion of *equilibrium in beliefs* as defined in [37] when restricted to players having CPT preferences.¹²

In the setting of an n -player normal form game with real valued payoff functions, the pure Nash equilibria do not depend on the specific CPT features of the players, i.e. the reference point, the value function and the two probability weighting functions, one for gains and one for losses. Hence the traditional result on the lack of guarantee for the existence of a pure Nash equilibrium continues to hold when players have CPT preferences. Keskin [74] proves the existence of a mixed action Nash equilibrium for any finite game when players have CPT preferences. In Example 4.3.9, we show that a finite game may not have any black-box strategy Nash equilibrium.¹³ On the other hand, in Theorem 4.3.12, we prove our main result that for any finite game with players having CPT preferences, there exists a mixed black-box strategy Nash equilibrium. If the players have EUT preferences, then the notions of black-box strategy Nash equilibrium and mixed black-box strategy Nash equilibrium are equivalent to the notion of mixed action Nash equilibrium (when interpreted appropriately; see the remark before Proposition 4.3.14; see also Figure 4.5).

The chapter is organized as follows. In Section 4.2, we describe the CPT setup and establish that under this setup betweenness is equivalent to independence (Theorem 4.2.3).¹⁴ In Section 4.3, we describe an n -player non-cooperative game setup and define various notions of Nash equilibrium in the absence of betweenness, in particular with CPT preferences. We discuss the questions concerning their existence and how these different notions of equilibria compare with each other. In Section 5.6, we conclude with a table that summarizes the results.

4.2 CPT and Betweenness

In this chapter, we assume that each person is associated with a fixed *reference point* $r \in \mathbb{R}$, a *value function* $v : \mathbb{R} \rightarrow \mathbb{R}$, and two *probability weighting functions* $w^\pm : [0, 1] \rightarrow [0, 1]$, w^+ for gains and w^- for losses. The CPT value is evaluated as described in Section 1.3.

We now define some axioms for preferences over lotteries. We are interested in “mixtures” of lotteries, i.e. lotteries with other lotteries as outcomes. Consider a (two-stage) compound lottery $K := \{(q^j, L^j)\}_{1 \leq j \leq t}$, where $L^j = (p^j, z^j)$, $1 \leq j \leq t$, are lotteries over real outcomes and q^j is the chance of lottery L^j . We assume that $\sum_{j=1}^t q^j = 1$. A two-stage compound lottery can be reduced to a single-stage lottery by multiplying the probability vector p^j corresponding to the lottery L^j by q^j for each j , $1 \leq j \leq t$, and then adding the probabilities of identical outcomes across all the lotteries L^j , $1 \leq j \leq t$. Let $\sum_{j=1}^t q^j L^j$ denote the reduced lottery corresponding to the compound lottery K .

Let \preceq denote a preference relation over single-stage lotteries. We assume \preceq to be a weak order, i.e. \preceq is transitive (if $L_1 \preceq L_2$ and $L_2 \preceq L_3$, then $L_1 \preceq L_3$) and complete (for all L_1, L_2 , we have $L_1 \preceq L_2$ or $L_2 \preceq L_1$, where possibly both preferences hold). The additional binary relations \succeq, \sim, \prec and \succ are derived from \preceq in the usual manner. A preference relation \preceq is a CPT preference relation if there exist CPT features (r, v, w^\pm) such that $L_1 \preceq L_2$ iff $V(L_1) \leq V(L_2)$. Note that a CPT preference relation is a weak order. A preference relation \preceq satisfies *independence* if for any lotteries L_1, L_2 and L , and any constant $0 \leq \alpha \leq 1$, $L_1 \preceq L_2$ implies $\alpha L_1 + (1 - \alpha)L \preceq \alpha L_2 + (1 - \alpha)L$. A preference relation \preceq satisfies *betweenness* if for any lotteries $L_1 \preceq L_2$, we have $L_1 \preceq \alpha L_1 + (1 - \alpha)L_2 \preceq L_2$, for all $0 \leq \alpha \leq 1$. A preference relation \preceq satisfies *weak betweenness* if for any lotteries $L_1 \sim L_2$, we have $L_1 \sim \alpha L_1 + (1 - \alpha)L_2$, for all $0 \leq \alpha \leq 1$.

Suppose a preference relation \preceq satisfies independence. Then $L_1 \preceq L_2$ implies

$$L_1 = \alpha L_1 + (1 - \alpha)L_1 \preceq \alpha L_1 + (1 - \alpha)L_2 \preceq \alpha L_2 + (1 - \alpha)L_2 = L_2.$$

Thus, if a preference relation satisfies independence, then it satisfies betweenness. Also, if a preference relation satisfies betweenness, then it satisfies weak betweenness.

In the following example, we will provide CPT features for Alice so that her preferences agree with those described in Section 4.1. This example also shows that cumulative prospect theory can give rise to preferences that do not satisfy betweenness.

Example 4.2.1. Recall that Alice is faced with the following three lotteries:

$$\begin{aligned} L_1 &= \{(0.34, \$20,000); (0.66, \$0)\}, \\ L_2 &= \{(0.17, \$30,000); (0.83, \$0)\}, \\ L &= \{(0.01, \$30,000); (0.32, \$20,000); (0.67, \$0)\}. \end{aligned}$$

Let $r = 0$ be the reference point of Alice. Thus all the outcomes lie in the gains domain. Let $v(x) = x^{0.8}$ for $x \geq 0$; Alice is risk-averse in the gains domain. Let the probability weighting function for gains be given by

$$w^+(p) = \exp\{-(-\ln p)^{0.6}\},$$

a form suggested by Prelec [113] (see Figure 4.1). We won't need the probability weighting function for losses. Direct computations show that $V(L_1) = 968.96$, $V(L_2) = 932.29$, and




Figure 4.1: The solid curve shows the probability weighting function for Alice from Example 4.2.1 and Example 4.3.7, and the dashed curve shows the probability weighting function for Charlie from Example 4.2.2.

$V(L) = 1022.51$ (all decimal numbers in this example are correct to two decimal places). Thus the preference behavior of Alice, as described in Section 4.1 (i.e., she prefers L_1 over L_2 , but prefers L over L_1 and L_2), is consistent with CPT and can be modeled, for example, with the CPT features stated here. \square

The following example shows that CPT can give rise to preferences that do not satisfy weak betweenness (the lotteries and the CPT features presented below appear in [74]).

Example 4.2.2. Suppose Charlie has $r = 0$ as his reference point and $v(x) = x$ as his value function. Let his probability weighting function for gains be given by

$$w^+(p) = \exp\{-(-\ln p)^{0.5}\}.$$

(See Figure 4.1.) We won't need the probability weighting function for losses since we consider only outcomes in the gains domain in this example. Consider the lotteries $L_1 = \{(0.5, 2\beta); (0.5, 0)\}$ and $L_2 = \{(0.5, \beta + 1); (0.5, 1)\}$, where $\beta = 1/w^+(0.5) = 2.299$ (all decimal numbers in this example are correct to three decimal places). Direct computations reveal that $V(L_1) = V(L_2) = 2.000 > V(0.5L_1 + 0.5L_2) = 1.985$. \square

Given a utility function $u : \mathbb{R} \rightarrow \mathbb{R}$ (assumed to be continuous and strictly increasing), the expected utility of a lottery $L = \{(p_k, z_k)\}_{1 \leq k \leq m}$ is defined as $U(L) := \sum_{k=1}^m p_k u(z_k)$. A preference relation \preceq is said to be an EUT preference relation if there exists a utility function u such that $L_1 \preceq L_2$ iff $U(L_1) \leq U(L_2)$. Note that if the CPT probability weighting

functions are linear, i.e. $w^\pm(p) = p$ for $0 \leq p \leq 1$, then the CPT value of a lottery coincides with the expected utility of that lottery with respect to the utility function $u = v$. It is well known that EUT preference relations satisfy independence and hence betweenness. Several generalizations of EUT have been obtained by weakening the independence axiom and assuming only betweenness, for example, weighted utility theory [32, 33], skew-symmetric bilinear utility [47, 18], implicit expected utility [40, 34] and disappointment aversion theory [55, 17]. The following theorem shows that in the restricted setting of CPT preferences, betweenness and independence are equivalent.

Theorem 4.2.3. *If \preceq is a CPT preference relation, then the following are equivalent:*

- (i) \preceq is an EUT preference relation,
- (ii) \preceq satisfies independence,
- (iii) \preceq satisfies betweenness.

Wakker [131] considers rank-dependent utility (RDU) preferences [114], which is a special case of CPT preferences; RDU preferences are CPT preferences for which the probability weighting functions satisfy $w^+(p) = 1 - w^-(1 - p)$ for all $p \in [0, 1]$. Wakker proves that, under RDU preferences \preceq , the probability weighting function w^+ is linear if and only if \preceq satisfies betweenness. In fact, Wakker proves something more general. A preference relation \preceq is said to be quasi-concave (resp. quasi-convex) if for any lotteries L_1, L_2 and any constant $0 \leq \alpha \leq 1$, $L_1 \preceq L_2$ implies $L_1 \preceq \alpha L_1 + (1 - \alpha)L_2$ (resp. $L_1 \succeq L_2$ implies $L_1 \succeq \alpha L_1 + (1 - \alpha)L_2$). Note that \preceq satisfies betweenness if it is both quasi-concave and quasi-convex. Wakker shows that, under RDU preferences, w^+ is concave (resp. convex) if and only if \preceq is quasi-concave (resp. quasi-convex).

Wakker proves this by defining a measure of convexity

$$\lambda[p, q] := \frac{w^+(p)/2 + w^+(q)/2 - w^+(p/2 + q/2)}{w^+(p) - w^+(q)},$$

and showing that $\lambda[p_1, q_2] + \lambda[p_2, q_2] \geq 0$, for any $0 \leq p_1 < q_1 \leq p_2 < q_2 \leq 1$, if \preceq satisfies quasi-convexity by consideration of proper lotteries. A simple analytic proof is then used to show that the above condition implies convexity of w^+ . Although Wakker's proof can be easily modified to account for general CPT preferences, we give an alternative proof in Appendix 4.B, where we show that the probability weighting functions satisfy

$$[w^\pm(a_2) - w^\pm(b)] [w^\pm(b) - w^\pm(c_1)] = [w^\pm(b) - w^\pm(a_1)] [w^\pm(c_2) - w^\pm(b)],$$

for any $0 \leq a_1 < c_1 < b < c_2 < a_2 \leq 1$ such that $(a_2 - b)(b - c_1) = (b - a_1)(c_2 - b)$, if \preceq satisfies betweenness. We then solve this functional equation using the appropriate boundary conditions and the continuity property of w^\pm to show that $w^\pm(p) = p$ for $p \in [0, 1]$.

4.3 Equilibrium in Black-Box Strategies

We now consider an n -player non-cooperative game where the players have CPT preferences. We will discuss several notions of equilibrium for such a game and will contrast them.

Let $\Gamma := (N, (A_i)_{i \in N}, (x_i)_{i \in N})$ denote a *game*, where $N := \{1, \dots, n\}$ is the set of *players*, A_i is the finite *action set* of player i , and $x_i : A \rightarrow \mathbb{R}$ is the *payoff function* for player i . Here $A := \prod_i A_i$ denotes the set of all *action profiles* $a := (a_1, \dots, a_n)$. Let $A_{-i} := \prod_{i \neq j} A_j$ denote the set of action profiles a_{-i} of all players except player i .

Definition 4.3.1. For any action profile $a_{-i} \in A_{-i}$ of the opponents, we define the *best response action set* of player i to be

$$\mathcal{A}_i(a_{-i}) := \arg \max_{a_i \in A_i} x_i(a_i, a_{-i}). \quad (4.3.1)$$

Definition 4.3.2. An action profile $a = (a_1, \dots, a_n)$ is said to be a *pure Nash equilibrium* if for each player $i \in N$, we have

$$a_i \in \mathcal{A}_i(a_{-i}).$$

The notion of pure Nash equilibrium is the same whether the players have CPT preferences or EUT preferences because only deterministic lotteries, comprised of being offered one outcome with probability 1, are considered in the framework of this notion. It is well known that for any given game Γ , a pure Nash equilibrium need not exist.

Let $\mu_{-i} \in \Delta(A_{-i})$ denote a *belief* of player i on the action profiles of her opponents. Given the belief μ_{-i} of player i , if she decides to play action a_i , then she will face the lottery $\{(\mu_{-i}[a_{-i}], x_i(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}}$.

Definition 4.3.3. For any belief $\mu_{-i} \in \Delta(A_{-i})$, define the *best response action set* of player i as

$$\mathcal{A}_i(\mu_{-i}) := \arg \max_{a_i \in A_i} V_i \left(\{(\mu_{-i}[a_{-i}], x_i(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}} \right). \quad (4.3.2)$$

Note that this definition is consistent with the definition of the best response action set that takes an action profile a_{-i} of the opponents as its input (Definition 4.3.1), if we interpret a_{-i} as the belief $\mathbf{1}\{a_{-i}\} \in \Delta(A_{-i})$, since $\mathcal{A}_i(\mathbf{1}\{a_{-i}\}) = \mathcal{A}_i(a_{-i})$.

Let $\sigma_i \in \Delta(A_i)$ denote a *conjecture* over the action of player i . Let $\sigma := (\sigma_1, \dots, \sigma_n)$ denote a *profile of conjectures*, and let $\sigma_{-i} := (\sigma_j)_{j \neq i}$ denote the profile of conjectures for all players except player i . Let $\mu_{-i}(\sigma_{-i}) \in \Delta(A_{-i})$ be the belief induced by conjectures $\sigma_j, j \neq i$, given by

$$\mu_{-i}(\sigma_{-i})[a_{-i}] := \prod_{j \neq i} \sigma_j[a_{-i}],$$

which is nothing but the product distribution induced by σ_{-i} .

Definition 4.3.4. A conjecture profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is said to be a *mixed action Nash equilibrium* if, for each player i , we have

$$a_i \in \mathcal{A}_i(\mu_{-i}(\sigma_{-i})), \text{ for all } a_i \in \text{supp } \sigma_i.$$

In other words, the conjecture σ_i over the action of player i should assign positive probabilities to only optimal actions of player i , given her belief $\mu_{-i}(\sigma_{-i})$.

It is well known that a mixed Nash equilibrium exists for every game with EUT players, see [99]. Keskin [74] generalizes the result of Nash [99] on the existence of a mixed action Nash equilibrium to the case when players have CPT preferences.

Let $B_i := \Delta(A_i)$ denote the set of all black-box strategies for player i with a typical element denoted by $b_i \in B_i$. Recall that if player i implements a black-box strategy b_i , then we interpret this as a trusted party other than the player sampling an action $a_i \in A_i$ from the distribution b_i and playing action a_i on behalf of player i . We assume the usual topology on B_i . Let $B := \prod_i B_i$ and $B_{-i} := \prod_{j \neq i} B_j$ with typical elements denoted by b and b_{-i} , respectively.

Note that, although a conjecture σ_i and a black-box strategy b_i are mathematically equivalent, viz. they are elements of the same set $B_i = \Delta(A_i)$, they have different interpretations. We will call $s_i \in \Delta(A_i)$ a *mixture* of actions of player i when we want to be agnostic to which interpretation is being imposed. Let $S_i := \Delta(A_i)$, $S := \prod_i \Delta(A_i)$ and $S_{-i} := \prod_{j \neq i} S_j$ with typical elements denoted by s_i , s and s_{-i} , respectively. (Note that $S \neq \Delta(A)$ unless all but one player have singleton action sets.)

For any belief $\mu_{-i} \in \Delta(A_{-i})$ and any black-box strategy b_i of player i , let $\mu(b_i, \mu_{-i}) \in \Delta(A)$ denote the product distribution given by

$$\mu(b_i, \mu_{-i})[a] := b_i[a_i] \mu_{-i}[a_{-i}].$$

Given the belief μ_{-i} of player i , if she decides to implement the black-box strategy b_i , then she will face the lottery $\{\mu(b_i, \mu_{-i})[a], x_i(a)\}_{a \in A}$.

Definition 4.3.5. For any belief $\mu_{-i} \in \Delta(A_{-i})$, define the *best response black-box strategy set* of player i as

$$\mathcal{B}_i(\mu_{-i}) := \arg \max_{b_i \in B_i} V_i(\{(\mu(b_i, \mu_{-i})[a], x_i(a))\}_{a \in A}).$$

Lemma 4.3.6. *For any belief μ_{-i} , the set $\mathcal{B}_i(\mu_{-i})$ is non-empty, and*

$$\overline{\text{co}}(\mathcal{B}_i(\mu_{-i})) = \text{co}(\mathcal{B}_i(\mu_{-i})).$$

See Appendix 4.C for proof.

Let us compare the two concepts: the best response action set (Definition 4.3.3) and the best response black-box strategy set (Definition 4.3.5). Even though both of them take the belief μ_{-i} of player i as input, the best response action set $\mathcal{A}_i(\mu_{-i})$ outputs a collection of actions of player i , whereas the best response black-box strategy set $\mathcal{B}_i(\mu_{-i})$ outputs a collection of black-box strategies of player i , which are probability distributions over the set of actions $a_i \in A_i$. If we interpret an action a_i as the mixture $\mathbf{1}\{a_i\} \in S_i = \Delta(A_i)$, and a black-box strategy b_i as a mixture as well, then we can compare the two sets $\mathcal{A}(\mu_{-i})$ and $\mathcal{B}(\mu_{-i})$ as subsets of S_i . The following example shows that, in general, the two sets can be disjoint, and hence quite distinct.

	1	2	3
1	\$20,000	\$20,000	\$0
2	\$30,000	\$0	\$0

Figure 4.1: Payoff matrix for Alice in Example 4.3.7. Rows and columns correspond to Alice's and Bob's actions respectively. The amount in each cell corresponds to Alice's payoff.

Example 4.3.7. We consider a 2-player game. Let Alice be player 1, with action set $A_1 = \{1, 2\}$, and let Bob be player 2, with action set $A_2 = \{1, 2, 3\}$. Let the payoff function for Alice be as shown in Figure 4.1. Let $\mu_{-1} = (0.17, 0.17, 0.66) \in \Delta(A_{-1}) = \Delta(A_2)$ be the belief of Alice. Then, as considered in Section 4.1, Alice faces the lottery $L_1 = \{(0.34, \$20,000); (0.66, \$0)\}$ if she plays action 1 and the lottery $L_2 = \{(0.17, \$30,000); (0.83, \$0)\}$ if she plays action 2. We retain the CPT features for Alice, as in Example 4.2.1, viz.: $r = 0$, $v(x) = x^{0.8}$ for $x \geq 0$, and

$$w^+(p) = \exp\{-(-\ln p)^{0.6}\}.$$

We saw that $V_1(L_1) = 968.96$, $V_1(L_2) = 932.29$, and $V(16/17L_1 + 1/17L_2) = 1022.51$ (all decimal numbers in this example are correct to two decimal places). Amongst all the mixtures, the maximum CPT value is achieved at the unique mixture $L^* = \alpha^*L_1 + (1 - \alpha^*)L_2$, where $\alpha^* = 0.96$; we have $V_1(L^*) = 1023.16$. Thus, $\mathcal{A}_1(\mu_{-1}) = \{\mathbf{1}\{1\}\}$ and $\mathcal{B}_1(\mu_{-1}) = \{(\alpha^*, 1 - \alpha^*)\}$. \square

For any black-box strategy profile b_{-i} of the opponents, let $\mu_{-i}(b_{-i}) \in \Delta(A_{-i})$ be the induced belief given by

$$\mu_{-i}(b_{-i})[a_{-i}] := \prod_{j \neq i} b_j[a_{-i}].$$

Definition 4.3.8. A black-box strategy profile $b = (b_1, \dots, b_n)$ is said to be a *black-box strategy Nash equilibrium* if, for each player i , we have

$$b_i \in \mathcal{B}_i(\mu_{-i}(b_{-i})).$$

If the players have EUT preferences, a conjecture profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is a mixed action Nash equilibrium if and only if the black-box strategy profile $b = (b_1, \dots, b_n)$, where $b_i = \sigma_i$, for all $i \in N$, is a black-box strategy Nash equilibrium. Thus, under EUT, the notion of a black-box strategy Nash equilibrium is equivalent to the notion of a mixed action Nash equilibrium, although there is still a conceptual difference between these two notions based on the interpretations for the mixtures of actions. Further, we have the existence of a black-box strategy Nash equilibrium for any game when players have EUT preferences from the well-known result about the existence of a mixed action Nash equilibrium. The following example shows that, in general, a black-box strategy Nash equilibrium may not exist when players have CPT preferences.

	0	1
0	4	0
1	3	1

	0	1
0	0	1
1	1	0

Figure 4.2: Payoff matrices for the 2×2 game in Example 4.3.9 (left matrix for player 1 and right matrix for player 2). The rows and the columns correspond to the actions of player 1 and player 2, respectively, and the entries in the cell represent the corresponding payoffs.

Example 4.3.9. Consider a 2×2 game (i.e a 2-player game where each player has two actions $\{0, 1\}$) with the payoff matrices as shown in Figure 4.2. Let the reference points be $r_1 = r_2 = 0$. Let $v_i(\cdot)$ be the identity function for $i = 1, 2$. Let the probability weighting functions for gains for the two players be given by

$$w_i^+(p) = \exp\{-(-\ln p)^{\gamma_i}\}, \text{ for } i = 1, 2,$$

where $\gamma_1 = 0.5$ and $\gamma_2 = 1$. We do not need the probability weighting functions for losses since all the outcomes lie in the gains domain for both the players. Notice that player 2 has EUT preferences since $w_2^+(p) = p$.

Suppose player 1 and player 2 play black-box strategies $(1 - p, p)$ and $(1 - q, q)$, respectively, where $p, q \in [0, 1]$. With an abuse of notation, we identify these black-box strategies by p and q , respectively. The corresponding lottery faced by player 1 is given by

$$L_1(p, q) := \{(\mu[0, 0], 4); (\mu[1, 0], 3); (\mu[1, 1], 1); (\mu[0, 1], 0)\},$$

where $\mu[0, 0] := (1 - p)(1 - q)$, $\mu[1, 0] := p(1 - q)$, $\mu[0, 1] := (1 - p)q$, and $\mu[1, 1] := pq$. By (1.3.8), the CPT value of the lottery faced by player 1 is given by

$$\begin{aligned} V_1(L_1(p, q)) &:= 4 \times [w_1^+(\mu[0, 0])] \\ &\quad + 3 \times [w_1^+(\mu[0, 0] + \mu[1, 0]) - w_1^+(\mu[0, 0])] \\ &\quad + 1 \times [w_1^+(\mu[0, 0] + \mu[1, 0] + \mu[1, 1]) - w_1^+(\mu[0, 0] + \mu[1, 0])]. \end{aligned}$$

The plot of the function $V_1(L_1(p, q))$ with respect to p , for $q = 0.3$ and $q = 0.35$, is shown in Figure 4.3. We observe that the best response black-box strategy set $\mathcal{B}_1(\mu_{-1}(q))$ of player 1 to player 2's black-box strategy $q \in B_2$ satisfies the following: $\mathcal{B}_1(\mu_{-1}(q)) = \{0\}$ for $q < q^*$, $\mathcal{B}_1(\mu_{-1}(q)) = \{0, p^*\}$ for $q = q^*$, and $\mathcal{B}_1(\mu_{-1}(q)) \subset [p^*, 1]$ for $q > q^*$, where $p^* = 0.996$ and $q^* = 0.340$ (here the numbers are correct to three decimal points). Further, $\mathcal{B}_1(\mu_{-1}(q))$ is singleton for $q \in (q^*, 1]$ and the unique element in $\mathcal{B}_1(\mu_{-1}(q))$ increases monotonically with respect to q from p^* to 1 (see Figure 4.4). In particular, $\mathcal{B}_1(\mu_{-1}(1)) = \{1\}$. The lottery faced by player 2 is given by

$$L_2(p, q) := \{(\mu[0, 0], 0); (\mu[1, 0], 1); (\mu[1, 1], 0); (\mu[0, 1], 1)\},$$




Figure 4.3: The CPT value of player 1 in Example 4.3.9. Here, p and q denote the black-box strategies for player 1 and 2, respectively. Note the rise and sharp drop in the two curves near $p = 1$. For the curve for $q = 0.3$, the global maximum is attained at $p = 0$, whereas, for the curve for $q = 0.35$, the global maximum is attained close to $p = 1$, specifically for some $p \in [0.9, 1]$.

and the CPT value of player 2 for this lottery is given by $V_2(L_2(p, q)) = p(1 - q) + q(1 - p)$. The best response black-box strategy set $\mathcal{B}_2(\mu_{-2}(p))$ of player 2 to player 1's black-box strategy $p \in B_1$ satisfies the following: $\mathcal{B}_2(\mu_{-2}(p)) = \{1\}$ for $p < 0.5$, $\mathcal{B}_2(\mu_{-2}(p)) = [0, 1]$ for $p = 0.5$, and $\mathcal{B}_2(\mu_{-2}(p)) = \{0\}$ for $p > 0.5$. As a result, see Figure 4.4, there does not exist any (p', q') such that $p' \in \mathcal{B}_1(\mu_{-1}(q'))$ and $q' \in \mathcal{B}_2(\mu_{-2}(p'))$, and hence no black-box strategy Nash equilibrium exists for this game. \square

Let $\tau_i \in \mathcal{P}(B_i)$ denote a conjecture over the black-box strategy of player i . This will induce a conjecture $\sigma_i(\tau_i) \in \Delta(A_i)$ over the action of player i , given by

$$\sigma_i(\tau_i)[a_i] = \mathbb{E}_{\tau_i} b_i[a_i].$$

Given conjectures over black-box strategies $(\tau_j \in \Delta(B_j), j \neq i)$, let $\sigma_{-i}(\tau_{-i}) := (\sigma_j(\tau_j))_{j \neq i}$.

Definition 4.3.10. A profile of conjectures over black-box strategies $\tau = (\tau_1, \dots, \tau_n)$ is said to be a *mixed black-box strategy Nash equilibrium* if, for each player i , we have

$$b_i \in \mathcal{B}_i(\mu_{-i}(\sigma_{-i}(\tau_{-i}))), \text{ for all } b_i \in \text{supp } \tau_i.$$

Proposition 4.3.11. For a profile of conjectures $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$, consider the condition

$$\sigma_i^* \in \overline{\text{co}}(\mathcal{B}_i(\mu_{-i}(\sigma_{-i}^*))), \text{ for all } i. \quad (4.3.3)$$

- (i) If τ is a mixed black-box strategy Nash equilibrium, then the profile of conjectures σ^* , where $\sigma_i^* = \sigma_i(\tau_i), \forall i$, satisfies (4.3.3).




Figure 4.4: The figure (not to scale) shows the best response black-box strategy sets of the two players for the game in Example 4.3.9. The red (dashed) line shows the best response black-box strategy set of player 2 in response to the black-box strategy $(1 - p, p)$ of player 1. The green (solid) line shows the best response black-box strategy set of player 1 in response to the black-box strategy $(1 - q, q)$ of player 2. Note that there is no intersection of these lines.

(ii) If σ^* satisfies (4.3.3), then there exists a profile of finite support conjectures on black-box strategies $\hat{\tau} = (\hat{\tau}_1, \dots, \hat{\tau}_n)$, where $\hat{\tau}_i \in \Delta_f(B_i), \forall i$, that is a mixed black-box strategy Nash equilibrium, such that $\sigma_i^* = \sigma_i(\hat{\tau}_i), \forall i$.

We prove this proposition in Appendix 4.D. The content of this proposition is that in order to determine whether a profile τ of conjectures on black box strategies is a mixed black-box strategy Nash equilibrium or not it suffices to study the associated profile of conjectures on actions that is induced by τ . This justifies the study of the set mBBNE discussed below.

Theorem 4.3.12. *For any game Γ , there exists a profile of conjectures $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ that satisfies (4.3.3).*

We prove this theorem in Appendix 4.E. We have the following corollary from Theorem 4.3.12 and statement (ii) of Proposition 4.3.11.

Corollary 4.3.13. *For any finite game Γ , there exists a mixed black-box strategy Nash equilibrium. In particular, there is one that is a profile of finite support conjectures over the black-box strategies of players.* \square

We now compare the different notions of Nash equilibrium defined above. To that end, we will associate each of the equilibrium notions with their corresponding natural profile of mixtures over actions. For example, corresponding to any pure Nash equilibrium $a =$

(a_1, \dots, a_n) , assign the profile of mixtures over actions $(\mathbf{1}\{a_1\}, \dots, \mathbf{1}\{a_n\}) \in S$. Let $\text{pNE} \subset S$ denote the set of all profiles of mixtures over actions that correspond to pure Nash equilibria. Let $\text{mNE} \subset S$ denote the set of all mixed action Nash equilibria $\sigma \in S$. Let $\text{BBNE} \subset S$ denote the set of all black-box strategy Nash equilibria $b \in S$. Corresponding to any mixed black-box strategy Nash equilibrium $\tau = (\tau_1, \dots, \tau_n)$, assign the profiles of mixtures over actions $(\sigma_1(\tau_1), \dots, \sigma_n(\tau_n)) \in S$, and let $\text{mBBNE} \subset S$ denote the set of all such profiles. Note that each of the above subsets depends on the underlying game Γ and the CPT features of the players.

Proposition 4.3.14. *For any fixed game Γ and CPT features of the players, we have*

(i) $\text{pNE} \subset \text{mNE}$,

(ii) $\text{pNE} \subset \text{BBNE}$, and

(iii) $\text{BBNE} \subset \text{mBBNE}$.

Proof. The proof of statement (i) can be found in [74].

For statement (ii), let $(\mathbf{1}\{a_1\}, \dots, \mathbf{1}\{a_n\}) \in \text{pNE}$. For a black-box strategy b_i of player i , the belief $\mu_{-i} = \mathbf{1}\{a_{-i}\}$ of player i gives rise to the lottery $\{(b_i[a'_i], x_i(a'_i, a_{-i}))\}_{a'_i \in A_i}$. From the definition of CPT value (see Equation (1.3.8)), we observe that $V_i(\{(b_i[a'_i], x_i(a'_i, a_{-i}))\}_{a'_i \in A_i})$ is optimal as long as the probability distribution b_i does not assign positive probability to any suboptimal outcome. Hence,

$$\mathcal{B}_i(\mathbf{1}\{a_{-i}\}) = \text{co}(\mathbf{1}\{a'_i\} \in S_i : a'_i \in \mathcal{A}_i(\mathbf{1}\{a_{-i}\})).$$

In particular, $\mathbf{1}\{a_i\} \in \mathcal{B}_i(\mathbf{1}\{a_{-i}\})$, and hence $(\mathbf{1}\{a_1\}, \dots, \mathbf{1}\{a_n\}) \in \text{BBNE}$.

Statement (iii) follows directly from the Definitions 4.3.8 and 4.3.10. \square

In the following, we show via examples that each of the labeled regions ((a)–(g)), in Figure 4.5, is non-empty in general.

Example 4.3.15. For each of the seven regions in Figure 4.5, we provide a 2×2 game with the accompanying CPT features for the two players verifying that the corresponding region is non-empty. Let the action sets be $A_1 = A_2 = \{0, 1\}$. With an abuse of notation, let $p, q \in [0, 1]$ denote the mixtures over actions for players 1 and 2, respectively, where p and q are the probabilities corresponding to action 1 for both the players. Thus, the set of all profiles of mixtures over actions is $S = \{(p, q) : p, q \in [0, 1]\}$. Let $L_1(p, q)$ and $L_2(p, q)$ denote the corresponding lotteries faced by the two players. (All decimal numbers in these examples are correct to three decimal places.)

(a) Let both the players have EUT preferences with their utility functions given by the identity functions $u_i(x) = x$, for $i = 1, 2$. Let the payoff matrix be as shown in Figure 4.6a. Clearly, $(p = 0, q = 0) \in \text{pNE}$.

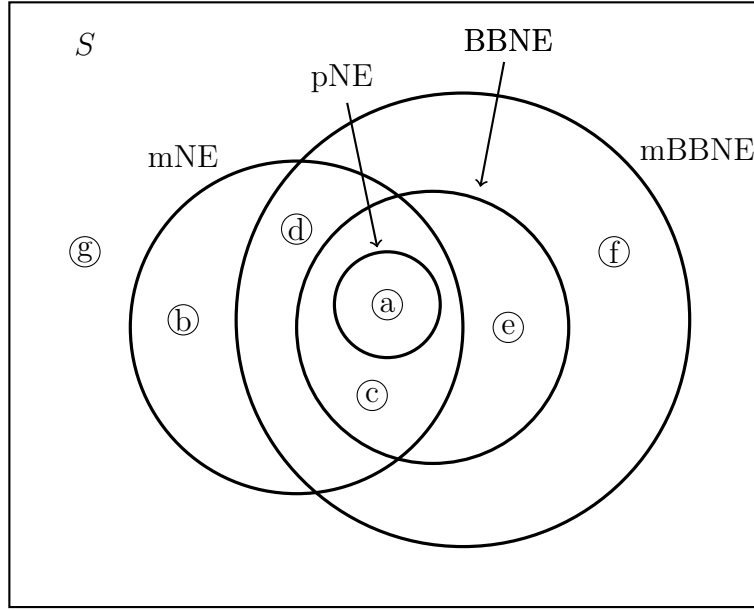


Figure 4.5: Venn diagram depicting the different notions of equilibrium as subsets of the set $S = \prod_i \Delta(A_i)$. The sets marked pNE, mNE, BBNE, and mBBNE represent the sets of pure Nash equilibria, mixed action Nash equilibria, black-box strategy Nash equilibria, and mixed black-box strategy Nash equilibria, respectively. Examples are given in the body of the text of CPT games lying in each of the indicated regions (a) through (g).

- (b) Let $r_i = 0, v_i(x) = x$, for $i = 1, 2$. Let $w_1^+(p) = p^{0.5}$ and $w_2^+(p) = p$, for $p \in [0, 1]$. Let the payoff matrix be as shown in Figure 4.6b, where $\beta := 1/w_1^+(0.5) = 1.414$. We have

$$L_1(p, q) = \{((1-p)(1-q), 2\beta); (p(1-q), \beta+1); (pq, 1); ((1-p)q, 0)\}.$$

The way β is defined, we get $V_1(L_1(0, 0.5)) = V_1(L_1(1, 0.5)) = 2$. Also, observe that $V_2(L_2(0.5, 0)) = V_2(L_2(0.5, q)) = V_2(L_2(0.5, 1)), \forall q \in [0, 1]$. With these observations, we get that $(0.5, 0.5) \in \text{mNE}$. We have, $\arg \max_{p \in [0, 1]} V_1(L_1(p, 0.5)) = \{p'\}$, where $p' = 0.707$ (see Figure 4.7). Hence $0.5 \notin \text{co}(\mathcal{B}_1(\mu_{-1}(0.5)))$ and $(0.5, 0.5) \notin \text{mBBNE}$.

- (c) Let the CPT features for both the players be as in (b). Let the payoff matrix be as shown in Figure 4.6c, where $\beta := 1/w_1^+(0.5) = 1.414$ and $\gamma = (1-p')/p'$ (here $p' = 0.707$ as in (b)). As observed in (b), $\mathcal{B}_1(\mu_{-1}(0.5)) = \{p'\}$. From the definition of γ , we see that player 2 is indifferent between her two actions, given her belief p' over player 1's actions. Thus $(p', 0.5) \in (\text{mNE} \cap \text{BBNE}) \setminus \text{pNE}$.
- (d) Let $r_i = 0, v_i(x) = x$, for $i = 1, 2$. Let $w_1^-(p) = p^{0.5}, w_2^+(p) = p$. Let the payoff matrix be as shown in Figure 4.6d, where $\beta := 1/w_1^-(0.5) = 1.414$. Note that the payoffs for player 1 are negations of her payoffs in (b), and her probability weighing

function for losses is same as her probability weighing function for gains in (b). Thus her CPT value function $V_1(L_1(p, q))$ is the negation of her CPT value function in (b). In particular, we have $V_1(L_1(0, 0.5)) = V_1(L_1(1, 0.5)) > V_1(L_1(p, 0.5))$ for all $p \in (0, 1)$. Thus, $0.5 \in \text{co}(\mathcal{B}_1(\mu_{-1}(0.5)))$, but $0.5 \notin \mathcal{B}_1(\mu_{-1}(0.5))$. The payoffs and CPT features of player 2 are same as in (b). Thus, $(0.5, 0.5) \in (\text{mNE} \cap \text{mBBNE}) \setminus \text{BBNE}$.

- (e) Let the CPT features for both the players be as in (b). Let the payoff matrix be as shown in Figure 4.6e, where $\beta := 1/w_1^+(0.5) = 1.414$, $\epsilon = 0.1$, and $\gamma := (1 - \tilde{p})/\tilde{p}$; here $\tilde{p} = 0.582$ is the unique maximizer of $V_1(L_1(p, 0.5))$ (see Figure 4.8). We have $V_1(L_1(0, 0.5)) = 2.071 > 2 = V_1(L_1(1, 0.5))$ and $\arg \max_p V_1(L_1(p, 0.5)) = \{\tilde{p}\}$ with $V_1(L_1(\tilde{p}, 0.5)) = 2.125$. From the definition of γ , we see that player 2 is indifferent between her two actions, given her belief \tilde{p} over player 1's actions. Thus, $(\tilde{p}, 0.5) \in \text{BBNE} \setminus \text{mNE}$.
- (f) Let the CPT features be as in Example 4.3.9. Let $p^* = 0.996$ and $q^* = 0.340$ be the same as in Example 4.3.9. Let the payoff matrix be as shown in Figure 4.6f. Note that the payoffs for both the players are the same as in Example 4.3.9. Recall $\mathcal{B}_1(\mu_{-1}(q)) = 0$ for $q < q^*$, $\mathcal{B}_1(\mu_{-1}(q)) = \{0, p^*\}$ for $q = q^*$, and $\mathcal{B}_1(\mu_{-1}(q)) \subset [p^*, 1]$ for $q > q^*$, and hence $0.5 \in \text{co}(\mathcal{B}_1(\mu_{-1}(q^*)))$ and $0.5 \notin \mathcal{B}_1(\mu_{-1}(q^*))$. Further, from the definition of γ , we have $V_2(L_2(0.5, 0)) = V_2(L_2(0.5, q)) = V_2(L_2(0.5, 1))$, $\forall q \in [0, 1]$. Hence, $(0.5, q^*) \in \text{mBBNE} \setminus (\text{mNE} \cap \text{BBNE})$.
- (g) Finally, if we let the players have EUT preferences and the payoffs as in (a), then $(1, 0) \notin (\text{mNE} \cup \text{mBBNE})$.

□

4.4 Summary

In the study of non-cooperative game theory from a decision-theoretic viewpoint, it is important to distinguish between two types of randomization:

1. conscious randomizations implemented by the players, and
2. randomizations in conjectures resulting from the beliefs held by the other players about the behavior of a given player.

This difference becomes evident when the preferences of the players over lotteries do not satisfy betweenness, a weakened form of independence property. We considered n -player normal form games where players have CPT preferences, an important example of preference relation that does not satisfy betweenness. This gives rise to four types of Nash equilibrium notions, depending on the different types of randomizations. We defined these different notions of equilibrium and discussed the question of their existence. The results are summarized in Table 4.1.

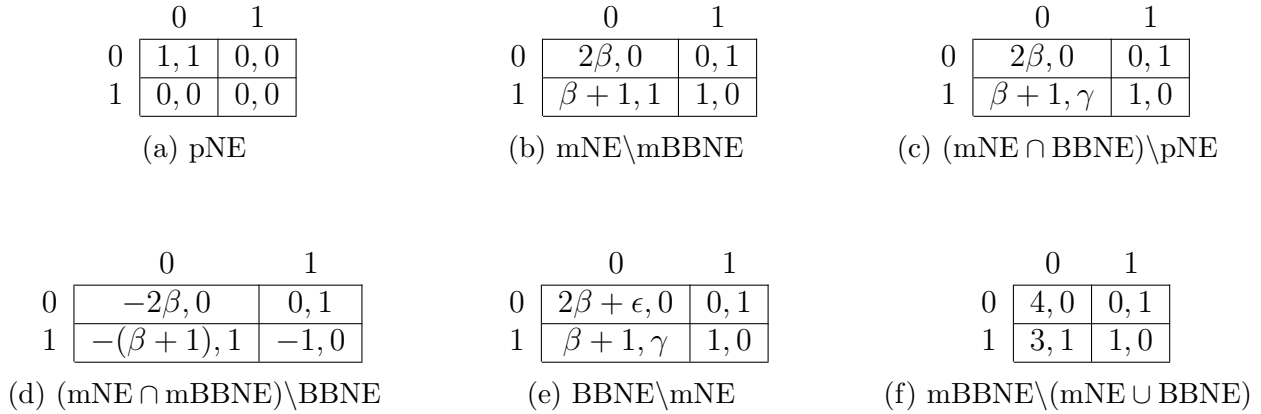


Figure 4.6: Payoff matrices for the 2×2 games in Example 4.3.15. The rows and the columns correspond to the actions of player 1 and player 2, respectively. In each cell, the left and right entries correspond to player 1 and player 2, respectively. The labels indicate the corresponding regions in Figure 4.5. The game matrix for the example corresponding to region (g) is the same as that for the one corresponding to region (a).




Figure 4.7: The CPT value function for player 1 in Example 4.3.15(b), when $q = 0.5$ is the mixture of actions of player 2.




Figure 4.8: The CPT value function for player 1 in Example 4.3.15(e), when $q = 0.5$ is the mixture of actions of player 2.

Type of Nash equilibrium	Strategies	Conjectures	Always exists
Pure Nash equilibrium	Pure actions	Exact conjectures	No
Mixed action Nash equilibrium	Pure actions	Mixed conjectures	Yes [74]
Black-box strategy Nash equilibrium	Black box strategies	Exact conjectures	No (Example 4.3.9)
Mixed black-box strategy Nash equilibrium	Black box strategies	Mixed conjectures	Yes (Theorem 4.3.12)

Table 4.1: Different types of Nash equilibrium when players have CPT preferences.

Appendix

4.A Proof of Theorem 4.2.3

Proof. Let the CPT preference relation \preceq be given by (r, v, w^\pm) . Since an EUT preference relation satisfies independence, we get that (i) implies (ii). Since betweenness is a weaker condition than independence, we get that (ii) implies (iii). We will now show that if \preceq satisfies betweenness, then the probability weighting functions are linear, i.e. $w^\pm(p) = p$ for $0 \leq p \leq 1$. This will imply that \preceq is an EUT preference relation with utility function $u = v$, and hence complete the proof.

Assume that the CPT preference relation \preceq satisfies betweenness. Consider a lottery $A := \{(p_1, z_1), (p_2, z_2), (1 - p_1 - p_2, r)\}$ such that $z_1 \geq z_2 \geq r$, $p_1 \geq 0$, $p_2 > 0$ and $p_1 + p_2 \leq 1$. By (1.3.9), we have

$$V(A) = \delta_1 w^+(P_1) + \delta_2 w^+(P_2),$$

where $\delta_1 := v(z_1) - v(z_2)$, $\delta_2 := v(z_2) - v(r)$, $P_1 := p_1$ and $P_2 := p_1 + p_2$. Let lottery $B := \{(q_1, z_1), (q_2, z_2), (1 - q_1 - q_2, r)\}$ be such that $q_1, q_2 \geq 0$, $Q_1 := q_1 > p_1$, and $Q_2 := q_1 + q_2 < P_2$. By (1.3.9), we have

$$V(B) = \delta_1 w^+(Q_1) + \delta_2 w^+(Q_2).$$

If z_1, z_2, p_1, p_2, q_1 and q_2 are such that

$$\frac{\delta_1}{\delta_2} = \frac{w^+(P_2) - w^+(Q_2)}{w^+(Q_1) - w^+(P_1)}, \quad (4.A.1)$$

then $V(A) = V(B)$ and, by betweenness, for any $0 \leq \alpha \leq 1$ we have $V(A) = V(B) = V(\alpha A + (1 - \alpha)B)$, i.e.

$$\delta_1 w^+(Q_1) + \delta_2 w^+(Q_2) = \delta_1 w^+(\alpha P_1 + (1 - \alpha)Q_1) + \delta_2 w^+(\alpha P_2 + (1 - \alpha)Q_2).$$

Using (4.A.1) we get

$$\begin{aligned} [w^+(P_2) - w^+(Q_2)] [w^+(Q_1) - w^+(\alpha P_1 + (1 - \alpha)Q_1)] \\ = [w^+(Q_1) - w^+(P_1)] [w^+(\alpha P_2 + (1 - \alpha)Q_2) - w^+(Q_2)]. \end{aligned} \quad (4.A.2)$$

Given any $0 \leq P_1 < Q_1 \leq Q_2 < P_2 \leq 1$, there exist z_1 and z_2 such that (4.A.1) holds. Indeed, take any $\delta > 0$ belonging to the range of the function v . This exists because $v(r) = 0$ and v is a strictly increasing function. Since w^+ is a strictly increasing function, we have

$$\kappa := \frac{w^+(P_2) - w^+(Q_2)}{w^+(Q_1) - w^+(P_1)} > 0.$$

Take $z_2 = v^{-1}(\delta/(1 + \kappa))$ and $z_1 = v^{-1}(\delta)$. These are well defined because v is assumed to be continuous and strictly increasing, and δ belongs to its range. Hence $z_1 > z_2 > r$ as required. Thus (4.A.2) holds for any $0 \leq P_1 < Q_1 \leq Q_2 < P_2 \leq 1$. In particular, when $Q_1 = Q_2$, we have

$$[w^+(P_2) - w^+(Q)] [w^+(Q) - w^+(R_1)] = [w^+(Q) - w^+(P_1)] [w^+(R_2) - w^+(Q)],$$

where $Q := Q_1 = Q_2$, $R_1 := \alpha P_1 + (1 - \alpha)Q$ and $R_2 := \alpha P_2 + (1 - \alpha)Q$. Equivalently, for any $0 \leq a_1 < c_1 < b < c_2 < a_2 \leq 1$ such that $(a_2 - b)(b - c_1) = (b - a_1)(c_2 - b)$, we have

$$[w^+(a_2) - w^+(b)] [w^+(b) - w^+(c_1)] = [w^+(b) - w^+(a_1)] [w^+(c_2) - w^+(b)].$$

In Lemma 4.B.1, we prove that the above condition implies $w^+(p) = p$, for $0 \leq p \leq 1$. Similarly, we can show that $w^-(p) = p$, for $0 \leq p \leq 1$. This completes the proof. \square

4.B An Interesting Functional Equation

Lemma 4.B.1. *Let $w : [0, 1] \rightarrow [0, 1]$ be a continuous, strictly increasing function such that $w(0) = 0$ and $w(1) = 1$. For any $0 \leq a_1 < c_1 < b < c_2 < a_2 \leq 1$ such that $(a_2 - b)(b - c_1) = (b - a_1)(c_2 - b)$, let*

$$[w(a_2) - w(b)] [w(b) - w(c_1)] = [w(b) - w(a_1)] [w(c_2) - w(b)]. \quad (4.B.1)$$

Then $w(p) = p$ for all $p \in [0, 1]$.

Proof. Taking $a_1 = 0, c_1 = 1/4, b = 1/2, c_2 = 3/4$ and $a_2 = 1$ in (4.B.1) we get,

$$[1 - w(1/2)] [w(1/2) - w(1/4)] = [w(1/2)] [w(3/4) - w(1/2)],$$

and hence,

$$w(3/4) = \frac{w(1/2) + w(1/2)w(1/4) - w(1/4)}{w(1/2)}.$$

Note that $w(1/2) > 0$. Taking $a_1 = 0, c_1 = 1/4, b = 1/3, c_2 = 1/2$ and $a_2 = 1$ in (4.B.1) we get,

$$[1 - w(1/3)] [w(1/3) - w(1/4)] = [w(1/3)] [w(1/2) - w(1/3)],$$

and hence,

$$w(1/3) = \frac{w(1/4)}{1 - w(1/2) + w(1/4)}.$$

Note that $1 - w(1/2) + w(1/4) > 1 - w(1/2) > 0$. Taking $a_1 = 0, c_1 = 1/3, b = 1/2, c_2 = 2/3$ and $a_2 = 1$ in (4.B.1) we get,

$$[1 - w(1/2)] [w(1/2) - w(1/3)] = [w(1/2)] [w(2/3) - w(1/2)],$$

and substituting for $w(1/3)$ we get,

$$w(2/3) = \frac{w(1/2) - w(1/2)^2 + 2w(1/2)w(1/4) - w(1/4)}{w(1/2) - w(1/2)^2 + w(1/2)w(1/4)}.$$

Note that

$$w(1/2) - w(1/2)^2 + w(1/2)w(1/4) = w(1/2)[1 - w(1/2) + w(1/4)] > 0.$$

Taking $a_1 = 0, c_1 = 1/2, b = 2/3, c_2 = 3/4$ and $a_2 = 1$ in (4.B.1) we get,

$$[1 - w(2/3)] [w(2/3) - w(1/2)] = [w(2/3)] [w(3/4) - w(2/3)].$$

Simplifying we get,

$$w(2/3) - w(2/3)w(3/4) = w(1/2) - w(1/2)w(2/3),$$

Substituting for $w(2/3)$ and $w(3/4)$ we get,

$$\begin{aligned} & \left[\frac{w(1/2) - w(1/2)^2 + 2w(1/2)w(1/4) - w(1/4)}{w(1/2) - w(1/2)^2 + w(1/2)w(1/4)} \right] \left[\frac{w(1/4) - w(1/2)w(1/4)}{w(1/2)} \right] \\ &= w(1/2) \left[\frac{w(1/4) - w(1/2)w(1/4)}{w(1/2) - w(1/2)^2 + w(1/2)w(1/4)} \right]. \end{aligned}$$

Since $w(1/4) - w(1/2)w(1/4) > 0$ and $w(1/2) - w(1/2)^2 + w(1/2)w(1/4) > 0$, we get

$$w(1/2) - w(1/4) = 2w(1/2)[w(1/2) - w(1/4)].$$

Since $w(1/2) - w(1/4) > 0$, we get $w(1/2) = 1/2$.

For any fixed $0 \leq x < y \leq 1$, let

$$w'(p') := \frac{w(p'(y-x) + x) - w(x)}{w(y) - w(x)}, \text{ for all } 0 \leq p' \leq 1.$$

Note that $w' : [0, 1] \rightarrow [0, 1]$ is a continuous, strictly increasing function with $w'(0) = 0$ and $w'(1) = 1$. Further, if $0 \leq a'_1 < c'_1 < b' < c'_2 < a'_2 \leq 1$ are such that $(a'_2 - b')(b' - c'_1) = (b' - a'_1)(c'_2 - b')$, then

$$[w'(a'_2) - w'(b')][w'(b') - w'(c'_1)] = [w'(b') - w'(a'_1)][w'(c'_2) - w'(b')].$$

Thus $w'(1/2) = 1/2$ and hence $w((x+y)/2) = (w(x) + w(y))/2$. Using this repeatedly we get $w(k/2^t) = k/2^t$, for $0 \leq k \leq 2^t$, $t = 1, 2, \dots$. Continuity of w then implies $w(p) = p$, for all $p \in [0, 1]$. \square

4.C Proof of Lemma 4.3.6

Proof. For a lottery $L = (p, z)$, where $z = (z_k)_{1 \leq k \leq m}$ is the outcome profile, and $(p_k)_{1 \leq k \leq m}$ is the probability vector, the function $V_i(p, z)$ is continuous with respect to $p \in \Delta^{m-1}$ [74]. Thus, $V_i(\{(\mu(b_i, \mu_{-i})[a], x_i(a))\}_{a \in A})$ is a function continuous with respect to $b_i \in B_i$, and hence $\mathcal{B}_i(\mu_{-i})$ is a non-empty closed subset of the compact space B_i . Since the convex hull of a compact subset of a Euclidean space is compact, the set $\text{co}(\mathcal{B}_i(\mu_{-i}))$ is closed. This completes the proof. \square

4.D Proof of Proposition 4.3.11

Proof. Suppose τ is a mixed black-box strategy Nash equilibrium. Let $\sigma_i^* = \sigma_i(\tau_i)$. Then, for all $b_i \in \text{supp } \tau_i$, we have $b_i \in \mathcal{B}_i(\mu_{-i}(\sigma_{-i}^*))$, and hence $\sigma_i^* \in \overline{\text{co}}(\mathcal{B}_i(\mu_{-i}(\sigma_{-i}^*)))$. This proves statement (i).

For statement (ii), suppose σ^* satisfies condition (4.3.3). In fact, by Lemma 4.3.6 we have, $\sigma_i^* \in \text{co}(\mathcal{B}_i(\mu_{-i}(\sigma_{-i}^*))) \subset \Delta(A_i)$, and by Caratheodory's theorem, σ_i^* is a convex combination of at most $|A_i|$ elements in $\mathcal{B}_i(\mu_{-i}(\sigma_{-i}^*))$. Hence, we can construct a mixed black-box strategy Nash equilibrium $\dot{\tau}$ such that $\dot{\tau}_i \in \Delta_f(B_i)$ and $\sigma_i^* = \sigma_i(\dot{\tau}_i), \forall i$. \square

4.E Proof of Theorem 4.3.12

Proof. The idea is to use the Kakutani fixed-point theorem, as in the proof of the existence of mixed action Nash equilibrium [100]. Assume the usual topology on S_i , for each i , and

let S have the corresponding product topology. The set S is a non-empty compact convex subset of the Euclidean space $\prod_i \mathbb{R}^{|A_i|}$. Let $K(\sigma)$ be the set-valued function given by

$$K(\sigma) := \prod_i \overline{\text{co}}(\mathcal{B}_i(\mu_{-i}(\sigma_{-i}))),$$

for all $\sigma \in S$. Since $\overline{\text{co}}(\mathcal{B}_i(\mu_{-i}(\sigma_{-i})))$ is non-empty and convex for each i (Lemma 4.3.6), the function $K(\sigma)$ is non-empty and convex for any $\sigma \in S$. We now show that the function $K(\cdot)$ has a closed graph. Let $\{\sigma^t\}_{t=1}^\infty$ and $\{\bar{\sigma}^t\}_{t=1}^\infty$ be two sequences in S that converge to $\bar{\sigma}$ and \bar{s} , respectively, and let $s^t \in K(\sigma^t)$ for all t . It is enough to show that $\bar{s} \in K(\bar{\sigma})$. For all $s_i \in S_i, \sigma_{-i} \in S_{-i}$, let

$$\tilde{V}_i(s_i, \sigma_{-i}) := \sup_{\substack{\tau_i \in \mathcal{P}(B_i), \\ \mathbb{E}_{\tau_i} b_i = s_i}} \mathbb{E}_{\tau_i} V_i(\{(\mu(b_i, \mu_{-i}(\sigma_{-i}))[a], x_i(a))\}_{a \in A}).$$

Since the product distribution $\mu(b_i, \mu_{-i}(\sigma_{-i}))$ is jointly continuous in b_i and σ_{-i} , and, as noted earlier, $V_i(p, z)$ is continuous with respect to the probability vector p , for any fixed outcome profile z , the function $V_i(\{(\mu(b_i, \mu_{-i}(\sigma_{-i}))[a], x_i(a))\}_{a \in A})$ is jointly continuous in b_i and σ_{-i} . This implies that the function $\tilde{V}_i(s_i, \sigma_{-i})$ is jointly continuous in s_i and σ_{-i} (see Appendix 4.F). From the definition of \tilde{V}_i , it follows that

$$\max_{s_i \in \Delta(A_i)} \tilde{V}_i(s_i, \sigma_{-i}) = \max_{b_i \in B_i} V_i(\{(\mu(b_i, \mu_{-i}(\sigma_{-i}))[a], x_i(a))\}_{a \in A}).$$

Indeed, the maximum on the left-hand side is well-defined since $\Delta(A_i)$ is a compact space and $\tilde{V}_i(\cdot, \sigma_{-i})$ is a continuous function. The maximum on the right-hand side is well-defined and the maximum is achieved by all $b_i \in \mathcal{B}_i(\mu_{-i}(\sigma_{-i}))$ (Lemma 4.3.6). Hence,

$$\arg \max_{s_i \in \Delta(A_i)} \tilde{V}_i(s_i, \sigma_{-i}) = \overline{\text{co}}(\mathcal{B}_i(\mu_{-i}(\sigma_{-i}))).$$

Since $s_i^t \in \overline{\text{co}}(\mathcal{B}_i(\mu_{-i}(\bar{\sigma}_{-i}^t)))$, for all t , we have

$$\tilde{V}_i(s_i^t, \sigma_{-i}^t) \geq \tilde{V}_i(\tilde{s}_i, \sigma_{-i}^t), \quad \text{for all } \tilde{s}_i \in S_i.$$

Since $\tilde{V}_i(s_i, \sigma_{-i})$ is jointly continuous in s_i and σ_{-i} , we get

$$\tilde{V}_i(\bar{s}_i, \bar{\sigma}_{-i}) \geq \tilde{V}_i(\tilde{s}_i, \bar{\sigma}_{-i}), \quad \text{for all } \tilde{s}_i \in S_i.$$

Hence we have $\bar{s}_i \in \overline{\text{co}}(\mathcal{B}_i(\mu_{-i}(\bar{\sigma}_{-i})))$. This shows that the function $K(\cdot)$ has a closed graph. By the Kakutani fixed-point theorem, there exists σ^* such that $\sigma^* \in K(\sigma^*)$, i.e. σ^* satisfies condition (4.3.3) [69]. This completes the proof. \square

4.F Joint Continuity of the Concave Hull of a Jointly Continuous Function

Let Δ^{m-1} and Δ^{n-1} be simplices of the corresponding dimensions with the usual topologies. Let $f : \Delta^{m-1} \times \Delta^{n-1} \rightarrow \mathbb{R}$ be a continuous function on $\Delta^{m-1} \times \Delta^{n-1}$ (with the product topology). Let $\mathcal{P}(\Delta^{m-1})$ denote the space of all probability measures on Δ^{m-1} with the topology of weak convergence. Let $g : \Delta^{m-1} \times \Delta^{n-1} \rightarrow \mathbb{R}$ be given by

$$g(x, y) := \sup \left\{ \mathbb{E}_{X \sim p} f(X, y) \mid p \in \mathcal{P}(\Delta^{m-1}), \mathbb{E}_{X \sim p} \text{id}(X) = x \right\}.$$

where $\text{id} : \Delta^{m-1} \rightarrow \Delta^{m-1}$ is the identity function $\text{id}(x) := x, \forall x \in \Delta^{m-1}$ and the expectation is over a random variable X taking values in Δ^{m-1} with the distribution p .

Proposition 4.F.1. *The function $g(x, y)$ is continuous on $\Delta^{m-1} \times \Delta^{n-1}$.*

Proof. We first prove that the function $g(x, y)$ is upper semi-continuous. Let $x_t \rightarrow x$ and $y_t \rightarrow y$. Let $\{g(x_{t_n}, y_{t_n})\}$ be a convergent subsequence of $\{g(x_t, y_t)\}$ with limit L . It is enough to show that the limit $L \leq g(x, y)$. Since for all n the set $\{p \in \mathcal{P}(\Delta^{m-1}), \mathbb{E}_{X \sim p} \text{id}(X) = x_{t_n}\}$ is compact, we know that there exists $p_{t_n} \in \mathcal{P}(\Delta^{m-1})$, such that $g(x_{t_n}, y_{t_n}) = \mathbb{E}_{X \sim p_{t_n}} [f(X, y_{t_n})]$ and $\mathbb{E}_{X \sim p_{t_n}} [\text{id}(X)] = x_{t_n}$. The sequence $\{p_{t_n}\}$ has a convergent subsequence, say $p_{t_{n_k}} \rightarrow \bar{p}$ (because $\mathcal{P}(\Delta^{m-1})$ is a compact space). Now, $\mathbb{E}_{X \sim \bar{p}} [\text{id}(X)] = \lim_k \mathbb{E}_{X \sim p_{t_{n_k}}} [\text{id}(X)] = \lim_k x_{t_{n_k}} = x$. Further, $\mathbb{E}_{X \sim p_{t_{n_k}}} [f(X, y_{t_{n_k}})] \rightarrow \mathbb{E}_{X \sim \bar{p}} [f(X, y)]$, since the product distributions $p_{t_{n_k}} \times \mathbf{1}\{y_{t_{n_k}}\}$, for all k , on $\Delta^{m-1} \times \Delta^{n-1}$, converge weakly to the product distribution $\bar{p} \times \mathbf{1}\{y\}$. Thus, $L = \mathbb{E}_{X \sim \bar{p}} [f(X, y)] \leq g(x, y)$ and the function $g(x, y)$ is upper-semicontinuous.

We now prove that the function $g(x, y)$ is lower semi-continuous. Let $x_t \rightarrow x$ and $y_t \rightarrow y$. The simplex Δ^{m-1} can be triangulated into finitely many other simplices, say T_1, \dots, T_k , whose vertices are x and some $m-1$ of the m vertices of Δ^{m-1} . Let (x_{t_n}) be any subsequence such that all $x_{t_n} \in T_j$ for some simplex. It is enough to show that the \liminf of the sequence $\{g(x_{t_n}, y_{t_n})\}$ is greater than or equal to $g(x, y)$. Let the other vertices of T_j be e_1, \dots, e_{m-1} . Let $z_{t_n} = (z_{t_n}^1, \dots, z_{t_n}^l)$ be the barycentric coordinates of x_{t_n} with respect to the simplex T_j , i.e.

$$x_{t_n} = (1 - z_{t_n}^1 - \dots - z_{t_n}^{m-1})x + z_{t_n}^1 e_1 + \dots + z_{t_n}^{m-1} e_{m-1}.$$

The function $g(x, y)$ is concave in x for any fixed y by construction. We have,

$$g(x_{t_n}, y_{t_n}) \geq (1 - z_{t_n}^1 - \dots - z_{t_n}^{m-1})g(x, y_{t_n}) + z_{t_n}^1 g(e_1, y_{t_n}) + \dots + z_{t_n}^{m-1} g(e_{m-1}, y_{t_n}).$$

Since $z_{t_n} \rightarrow (0, \dots, 0)$ and $g(e_1, y_{t_n}), \dots, g(e_{m-1}, y_{t_n})$ are all finite we get,

$$\liminf g(x_{t_n}, y_{t_n}) \geq \liminf g(x, y_{t_n}).$$

Let $\tilde{p} \in \mathcal{P}(\Delta^{m-1})$ be such that $\mathbb{E}_{X \sim \tilde{p}} [f(X, y)] = g(x, y)$ and $\mathbb{E}_{X \sim \tilde{p}} [\text{id}(X)] = x$. Then, $g(x, y_{t_n}) \geq \mathbb{E}_{X \sim \tilde{p}} [f(X, y_{t_n})]$, for all n , and hence,

$$\liminf g(x, y_{t_n}) \geq \liminf \mathbb{E}_{X \sim \tilde{p}} [f(X, y_{t_n})] = g(x, y).$$

This shows that the function $g(x, y)$ is lower semi-continuous.

Since the function $g(x, y)$ is upper and lower semi-continuous, it is continuous. \square

Notes

¹² Crawford [37] defines the notion of equilibrium in beliefs for 2-player games, but, as noted by Crawford, it can be easily extended to games with more than 2 players.

¹³ For the setting of 2-player games, this follows from the result of Crawford [37] on the existence of equilibrium in beliefs. A technical complex analysis argument is required to extend this result to more than 2 players. It appears that Crawford was aware of this (see footnote 9 in [37]). We provide an independent proof in Section 4.3 for the sake of completeness.

¹⁴ A similar result was proved by Wakker [131] in the setting of rank-dependent utility (RDU) preferences [114], which is a special case of CPT preferences. See the discussion following Theorem 4.2.3 for more on this.

Chapter 5

Mediated Correlated Equilibrium: Reconsidering CPT Correlated Equilibrium

5.1 Introduction

In the previous two chapters we took the neoclassical economics viewpoint of game theory that attempts to explain an equilibrium as a self-evident outcome of the optimal behavior of the participating players, assuming them to be rational. An alternative approach, called *learning in games*, is concerned with justifying equilibrium behavior via a dynamic process where the players learn from the past play and observations from the environment, and adapt accordingly [7, 52, 136]. In this chapter, we will be concerned with this alternative approach. Along the process, we will naturally come up with a modified version of the notion of CPT correlated equilibrium that will prove to be more appropriate in the settings of learning in games as well as mechanism design, which we consider in more detail in the next chapter.

It becomes even more important to consider non-EUT behavior in the theory of learning in games. For example, in a *repeated game*, Hart [59] argues that players tend to use simple procedures like *regret* minimization. A player i is said to have no regret¹⁵ if, for each pair of her actions a_i, \tilde{a}_i , she does not regret not having played action \tilde{a}_i whenever she played action a_i . Such regrets can simply be computed as the difference in the average payoffs received by the player from playing action \tilde{a}_i instead of action a_i , assuming the opponents stick to their actions. While evaluating such regrets in the real world, however, players who are modeled as evaluating lotteries according to CPT preferences are likely to exhibit different kinds of learning behavior than that exhibited by EUT players. The proposed model in this chapter is an attempt to handle these systematic deviations in learning, anticipated from the empirically observed behavioral features exhibited by human agents, as captured by CPT. We pose the following question: *How do the predictions of the theory of learning in games change if the players behave according to CPT?*

For a repeated game, Foster and Vohra [49] describe a procedure based on *calibrated learning* that guarantees the convergence of the *empirical distribution* of action play to the set of correlated equilibria, when players behave according to EUT. In Section 5.2, we formulate an analog for their procedure when players behave according to CPT. In Example 5.2.1, we describe a game for which the set of all CPT correlated equilibria is non-convex and we show that the empirical distribution of action play does not converge to this set.

We then define an extension of the set of CPT correlated equilibria and establish the convergence of the empirical distribution of action play to this extended set. It turns out that this extension has a nice game-theoretic interpretation, obtained by allowing the mediator to send any private signal (instead of restricting her to send a signal corresponding to some action). We formally define this setup in Section 5.3, and call it a *mediated game*. Myerson [91] has considered a further generalization in which each player i first reports her *type* from a finite set T_i . The mediator collects the reports from all the players and then sends each one of them a private signal from a finite set B_i . The mediator is characterized by a rule $\psi : \prod_i T_i \rightarrow \Delta(\prod_i B_i)$ that maps each type profile to a probability distribution on the set of signal profiles from which it samples the private signals to be sent. Based on her received signal, each player chooses her action. These are called *games with communication*. The type sets $(T_i)_{i=1}^n$, the signal sets $(B_i)_{i=1}^n$, and the mediator rule ψ together are said to comprise a *communication system*. Under EUT, the set of all correlated equilibria of a game is characterized as the union, over all possible communication systems, of the sets of joint distributions on the action profiles of all players arising from all the Nash equilibria for the corresponding game with communication (for a detailed exposition, see Chapter 6 in [89]). This is sometimes referred to as the *Bayes-Nash revelation principle*, or simply the *revelation principle*. Since a mediated game is a specific type of game with communication, characterized by players not reporting their type, or equivalently by the mediator ignoring the types reported by the players, our analysis shows that the revelation principle does not hold under CPT.

Calibrated learning is one way of studying learning in games. Some other approaches originate from Blackwell's *approachability* theory and the regret-based framework of online learning ([60, 53]). In fact, Foster and Vohra [49] establish the existence of calibrated learning schemes using such a regret-based framework and Blackwell's approachability theory. See [104] for a comparison between these approaches, and see also [27]. Hannan [56] introduced the concept of no-regret strategies in the context of repeated matrix games. No-regret learning in games is equivalent to the convergence of the empirical distribution of action play to the set of correlated equilibria [60, 53]. We establish an analog of this result when players behave according to CPT. We then ask if no-regret learning is possible under CPT.

Blackwell's approachability theorem prescribes a strategy to steer the average payoff vector of a player in a game with vector payoffs towards a given target set, irrespective of the strategies of the other players. The theorem also gives a necessary and sufficient condition for the existence of such a strategy provided the target set is convex and the game environment remains fixed. Here, by game environment, we mean the rule by which the payoff vectors depend on the players' actions. Under EUT, Hart and Mas-Colell [60] take these payoff

vectors to be the regrets associated to a player and establish no-regret learning by showing that the nonpositive orthant in the space of payoff vectors is approachable. Under CPT, although the target set is convex, the environment is not fixed. It depends on the empirical distribution of play at each step. A similar problem with dynamically evolving environment is considered in [70], where they get around this problem by considering a Stackelberg setting; one player (leader) plays an action first, then, after observing this action, the other player (follower) plays her action. In the absence of a Stackelberg setting, as in our case, we do not know of any result that characterizes approachability under dynamic environments. However, as far as games with CPT preferences are concerned, we answer this question by giving an example of a game for which a no-regret learning strategy does not exist (Example 5.5.2).

5.2 Calibrated Learning in Games

Let $\Gamma = ([n], (A_i)_{i \in [n]}, (x_i)_{i \in [n]})$ be a finite n -person game which is played repeatedly at each step $t \geq 1$. The game Γ is called the *stage game* of the repeated game. At every step t , each player i draws an action $a_i^t \in A_i$ with the probability distribution $\sigma_i^t \in \Delta(A_i)$. We assume that the randomizations of the players are independent of each other and of the past randomizations. For example, if each player i uses a uniform random variable U_i^t to draw a sample from σ_i^t , then the random variables $\{U_i^t\}_{i \in [n], t \geq 1}$ are independent. Each player is assumed to know the action space of all the players in the stage game Γ , but does not know the payoff functions and the CPT parameters of the other players. We assume that, after playing her action a_i^t , each player observes the actions taken by all the other players and thus at any step t all the players have access to the *past history* of the play at step t , $H^{t-1} := (a^1, \dots, a^{t-1})$, where $a^t := (a_i^t)_{i \in [n]}$ is the action profile played at step t . Let the strategy for player i for the repeated game above be given by $\tau_i := (\sigma_i^t, t \geq 1)$, where $\sigma_i^t : H^{t-1} \rightarrow \Delta(A_i)$, for each t .

We first describe the result of Foster and Vohra [50]. Suppose the players follow the following natural strategy: At every step t , on the basis of the past history of play, H^{t-1} , each player i predicts a joint distribution $\mu_{-i}^t \in \Delta(A_{-i})$ on the action profile of all the other players. This is player i 's *assessment* of how her opponents might play at step t . The sequence of functions of past history giving rise to the assessment is called the *assessment scheme* of the player. Depending on her assessment at step t , player i chooses a specific action among those that are most preferred by her in response to her assessment, called her *best reaction*.¹⁶ This is done using a fixed (time-invariant) function from $\Delta(A_{-i})$ to A_i , which maps $\mu_{-i} \in \Delta(A_{-i})$ to an action in A_i that is in the best response set for μ_{-i} ; this function is called the *best reaction map* of player i . Foster and Vohra [50] prove that (i) if each player's assessments are *calibrated* with respect to the sequence of action profiles of the other players and (ii) if each player plays the best reaction to her assessments, then the limit points of the empirical distribution of action play are correlated equilibria. By *action play* we mean the sequence of action profiles played by the players. We will give a formal definition of what is meant by calibration shortly. For the moment, roughly speaking, calibration says that the

empirical distributions conditioned on assessments converge to the assessments. The best reaction of player i to her assessment μ_{-i} of the actions of the other players, as considered in [50], is a specific action $a_i^* \in A_i$ that maximizes the expected payoff to player i with respect to her assessment, i.e.,

$$a_i^* \in \arg \max_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} \mu_{-i}(a_{-i}) x_i(a_i, a_{-i}).$$

Thus the best reaction is an action in the best response set. Note that it is assumed that each player uses a fixed tie breaking rule if there is more than one action in the best response set.

Suppose now that the players behave with CPT preferences. Given player i 's assessment μ_{-i} of the play of her opponents, she is faced with the following set of lotteries, one for each of her actions $a_i \in A_i$:

$$L_i(\mu_{-i}, a_i) := \{\mu_{-i}(a_{-i}), x_i(a_i, a_{-i})\}_{a_{-i} \in A_{-i}}.$$

Out of these lotteries, the ones she prefers most are those with the maximum CPT value $V_i(L_i(\mu_{-i}, a_i))$, evaluated using her CPT features. The choice of the action she takes corresponding to her most preferred lottery (with any arbitrary but fixed tie breaking rule) will be called her best reaction, and the map from $\Delta(A_{-i})$ to A_i giving the best reaction as a function of the assessment will be called the best reaction map of player i . Thus, once again, the best reaction is a specific action in the best response set.

We now ask the following question: *Suppose each player's assessments are calibrated with respect to the sequence of action profiles of the other players and she evaluates her best reaction in accordance with CPT preferences as explained above, then are the limit points of the empirical distribution of play contained in the set of CPT correlated equilibria?* Unfortunately, the answer is no (see Example 5.2.1). Before seeing why, let us give the promised formal definition of the notion of calibration.

Consider a sequence of outcomes y^1, y^2, \dots generated by Nature, belonging to some finite set S . At each step t , the forecaster predicts a distribution $q^t \in \Delta(S)$. Let $N(q, t)$ denote the number of times the distribution q is forecast up to step t , i.e. $N(q, t) := \sum_{\tau=1}^t \mathbf{1}\{q^\tau = q\}$, where $\mathbf{1}\{\cdot\}$ is the indicator function that takes value 1 if the expression inside $\{\cdot\}$ holds and 0 otherwise. Let $\rho(q, y, t)$ be the fraction of the steps on which the forecaster predicts q for which Nature plays $y \in S$, i.e.,

$$\rho(q, y, t) := \begin{cases} 0, & \text{if } N(q, t) = 0, \\ \frac{\sum_{\tau=1}^t \mathbf{1}\{q^\tau = q\} \mathbf{1}\{y^\tau = y\}}{N(q, t)}, & \text{otherwise.} \end{cases}$$

The forecast is said to be calibrated with respect to the sequence of plays made by Nature if

$$\lim_{t \rightarrow \infty} \sum_{q \in Q^t} |\rho(q, y, t) - q(y)| \frac{N(q, t)}{t} = 0, \text{ for all } y \in S, \quad (5.2.1)$$

where the sum is over the set Q^t of all distributions predicted by the forecaster up to step t .

	I	II	III	IV
0	$2\beta, 1$	$\beta + 1, 1$	$0, 1$	$1, 1$
1	$1.99, 0$	$1.99, 0$	$1.99, 0$	$1.99, 0$

Table 5.1: Payoff matrix for the game Γ^* in Example 5.2.1. The rows and columns correspond to player 1 and 2's actions respectively. The first entry in each cell corresponds to player 1's payoff and second to player 2's payoff.

Example 5.2.1. We consider a modification of the 2-player game proposed by Keskin [74], who uses it to demonstrate that the set of CPT correlated equilibria can be nonconvex. Let the 2-player game Γ^* be represented by the matrix in Table 5.1, where $\beta = 1/w_1^+(0.5)$. For the probability weighting functions $w_i^\pm(\cdot)$, we employ the functions of the form suggested by Prelec [113], which, for $i = 1, 2$, are given by

$$w_i^\pm(p) = \exp\{-(-\ln p)^{\gamma_i}\},$$

where $\gamma_1 = 0.5$ and $\gamma_2 = 1$. We thus have $w_1^+(0.5) = 0.435$ and $\beta = 2.299$. Let the reference points be $r_1 = r_2 = 0$. Let $v_i^{r_i}(\cdot)$ be the identity function for $i = 1, 2$. Notice that player 2 is indifferent amongst her actions.

Let $\mu_{odd} := (0.5, 0, 0.5, 0)$ and $\mu_{even} := (0, 0.5, 0, 0.5)$ be probability distributions on player 2's actions. We can evaluate the CPT values of player 1 for the following lotteries:

$$\begin{aligned} V_1(L_1(\mu_{odd}, 0)) &= 2\beta w_1^+(0.5) = 2, & V_1(L_1(\mu_{odd}, 1)) &= 1.99, \\ V_1(L_1(\mu_{even}, 0)) &= 1 + \beta w_1^+(0.5) = 2, & V_1(L_1(\mu_{even}, 1)) &= 1.99. \end{aligned}$$

Thus, player 1's best reaction to both these distributions μ_{odd} and μ_{even} is action 0. Since, player 2 is indifferent amongst her actions, we get that the distributions μ^o and μ^e , represented in Tables 5.2 and 5.3 respectively, belong to the set $C(\Gamma^*)$ (the set of CPT correlated equilibria of the game Γ^*). The mean of these two distributions is given by μ^* as represented in Table 5.4. Let $\mu_{unif} := (0.25, 0.25, 0.25, 0.25)$ be the uniform distribution on player 2's actions. The CPT values of player 1 for the lotteries corresponding to player 2 playing μ_{unif} are:

$$\begin{aligned} V_1(L_1(\mu_{unif}, 0)) &= w_1^+(0.75) + \beta w_1^+(0.5) + (\beta - 1)w_1^+(0.25) = 1.985, \\ V_1(L_1(\mu_{unif}, 1)) &= 1.99, \end{aligned}$$

since $w_1^+(0.25) = 0.308$ and $w_1^+(0.75) = 0.585$. We see that player 1's best reaction to the distribution μ_{unif} of player 2 is action 1. This shows that $\mu^* \notin C(\Gamma^*)$, and hence $C(\Gamma^*)$ is not convex.

Using this fact, we will attempt to construct an assessment scheme and a best reaction function for each player such that if each player makes assessments at each step according to her assessment scheme and acts according to the best reaction to her assessment at each

	I	II	III	IV
0	0.5	0	0.5	0
1	0	0	0	0

Table 5.2: Empirical distribution μ^o for the action play in Example 5.2.1.

	I	II	III	IV
0	0	0.5	0	0.5
1	0	0	0	0

Table 5.3: Empirical distribution μ^e for the action play in Example 5.2.1.

	I	II	III	IV
0	0.25	0.25	0.25	0.25
1	0	0	0	0

Table 5.4: Empirical distribution μ^* for the action play in Example 5.2.1.

step, then the assessments of each player are calibrated with respect to the sequence of action profiles of the other player and the limit of the generated empirical distribution of action play does not belong to $C(\Gamma^*)$.

Suppose player 2 plays her actions in a cyclic manner starting with action I at step 1, followed by actions II, III, IV and then again I and so on. Suppose player 1's assessment of player 2's action is $\mu_{odd} = (0.5, 0, 0.5, 0)$ and $\mu_{even} = (0, 0.5, 0, 0.5)$ at each odd and even step respectively. Then it is easy to see that player 1's assessments are calibrated with respect to the sequence of actions of player 2. (Here player 2 plays the role of Nature from the point of view of player 1.) Since player 1's best reaction is action 0 to all her assessments, she would play action 0 throughout. The distribution μ^* is a limit point of the empirical distribution of action play and does not belong to $C(\Gamma^*)$.

We have not described player 2's assessments. We would like to come up with an assessment scheme and a best reaction map for player 2 such that if player 2 forms assessments according to this assessment scheme and acts according to this best reaction map, then the sequence of her actions is the cyclic sequence that we require her to play and, further, player 2's assessments are calibrated with respect to the sequence of actions of player 1 (which is the all 0 sequence). Instead of doing this for the game Γ^* , we find it more natural to modify it into a 3-person game, denoted $\tilde{\Gamma}^*$, and create an assessment scheme and a best reaction map for each player in this 3-person game such that the assessments of each player are calibrated with respect to the sequence of action profiles of her opponents, each player plays her best reaction to her assessments at each step, and the limit empirical distribution of action play

exists but is not a CPT correlated equilibrium. We describe this in the following.

In the 3-person game $\tilde{\Gamma}^*$, player 1 has two actions $\{0,1\}$, and players 2 and 3 each have four actions $\{I,II,III,IV\}$. Let the payoffs to all the three players be -1 if players 2 and 3 play different actions. If players 2 and 3 play the same action, then let the resulting payoff matrix be as represented in Table 5.1, where the rows correspond to player 1's actions and the columns correspond to the common actions of players 2 and 3. Player 1 receives the payoff represented by the first entry in each cell and players 2 and 3 each receive the payoff represented by the second entry. Let player 1's CPT features be as in the 2-person game Γ^* . For players 2 and 3, let them be as for player 2 in that game. Let players 2 and 3 play in the cyclic manner as above, in sync with each other. Let player 1 play action 0 throughout. Let player 2's assessment at step t be the point distribution supported by the action profile a_{-2}^t which equals 0 for player 1 and the action played by player 2 for player 3. Similarly, let player 3's assessment at step t be the point distribution supported by the action profile a_{-3}^t which equals 0 for player 1 and the action played by player 3 for player 2. Then, for each of the players 2 and 3, her assessments are calibrated with respect to the sequence of action profiles of her opponents. Here the action pair comprised of the actions of players 1 and 3 plays the role of the actions of Nature from the point of view of player 2, and similarly the action pair comprised of the actions of players 1 and 2 plays the role of the actions of Nature from the point of view of player 3. The actions of player 2 and 3 at every step are best reactions to their corresponding assessments. Let the assessment of player 1 be $\tilde{\mu}_{odd}$ and $\tilde{\mu}_{even}$ at odd and even steps respectively, where now the distribution $\tilde{\mu}_{odd}$ puts 0.5 probability on action profiles (I,I) and (III,III), and $\tilde{\mu}_{even}$ puts 0.5 probability on action profiles (II,II) and (IV,IV). Again player 1's assessments are calibrated with respect to the sequence of action profiles of her opponents (where now action pairs comprised of the actions of player 2 and player 3 play the role of the actions of Nature from the point of view of player 1) and her actions are best reactions to her assessments. The limit point of the empirical distribution of action play is the distribution that puts probability 0.25 on action profiles (0,I,I), (0,II,II), (0,III,III) and (0,IV,IV). Since action 0 is not a best response of player 1 to the distribution $\tilde{\mu}_{unif}$ that puts probability 0.25 on action profiles (I,I), (II,II), (III,III) and (IV,IV), the limit point of the empirical distribution is not a CPT correlated equilibrium of the 3-player game $\tilde{\Gamma}^*$. Thus, we have a game where the assessments of each player are calibrated with respect to the sequence of action profiles of her opponents, each player plays her best reaction to her assessments at each step, and the limit empirical distribution of action play exists but is not a CPT correlated equilibrium. \square

5.3 Mediated Games and Equilibrium

In Example 5.2.1, the fact that action 0 is player 1's best reaction to the distributions μ_{odd} and μ_{even} , but not to μ_{unif} , plays an essential role in showing the non-convexity of the set $C(\Gamma^*)$ in the 2-player game Γ^* , and the fact that action 0 is player 1's best reaction to the distributions $\tilde{\mu}_{odd}$ and $\tilde{\mu}_{even}$, but not to $\tilde{\mu}_{unif}$, helps us in showing the non-convergence of

calibrated learning to the set $C(\tilde{\Gamma}^*)$ in the 3-player game $\tilde{\Gamma}^*$. We now describe a convex extension of the set $C(\Gamma)$ in a general finite n -person game Γ , and establish the convergence of the empirical distribution of action play to this extended set when each player plays the best reaction to her assessment at each step and her assessment scheme is calibrated with respect to the sequence of action profiles of her opponents. It turns out that this extended set of equilibria also has a game-theoretic interpretation, as follows. Suppose we add a signal system $(B_i)_{i \in [n]}$ to a game Γ , where each B_i is a finite set. (In Appendix 5.C, we study what happens when we relax the assumption that the sets B_i are finite and show that in a certain sense it is enough to consider only finite signal sets.) Suppose there is a mediator who sends a signal $b_i \in B_i$ to player i . Let $B := \prod_{i \in [n]} B_i$ be the set of all signal profiles $b = (b_i)_{i \in [n]}$, and let $B_{-i} := \prod_{j \neq i} B_j$ denote the set of signal profiles b_{-i} of all players except player i . Let $\tilde{\Gamma} := (\Gamma, (B_i)_{i \in [n]})$ denote such a game with a signal system. We call it a *mediated game*. The mediator is characterized by a distribution $\psi \in \Delta(B)$ that we call the *mediator distribution*. Thus, the mediator draws a signal profile $b = (b_i)_{i \in [n]}$ from the mediator distribution ψ and sends signal b_i to player i . Let ψ_i denote the marginal probability distribution on B_i induced by ψ , and for b_i such that $\psi_i(b_i) > 0$, let $\psi_{-i}(\cdot | b_i)$ denote the conditional probability distribution on B_{-i} . In the definition of a correlated equilibrium, the set B_i is restricted to be the set of actions A_i for each player i .

A *randomized strategy* for any player i is given by a function $\sigma_i : B_i \rightarrow \Delta(A_i)$ and a *randomized strategy profile* $\sigma = (\sigma_1, \dots, \sigma_n)$ gives the randomized strategy for all players. We define the *best response set* of player i to a randomized strategy profile σ and a mediator distribution ψ as

$$BR_i(\psi, \sigma) := \left\{ \sigma_i^* : B_i \rightarrow \Delta(A_i) \mid \text{for all } b_i \in \text{supp}(\psi_i), \right. \\ \left. \text{supp}(\sigma_i^*(b_i)) \subset \arg \max_{a_i \in A_i} V_i \left(\{ \tilde{\mu}_{-i}(a_{-i} | b_i), x_i(a_i, a_{-i}) \}_{a_{-i} \in A_{-i}} \right) \right\}, \quad (5.3.1)$$

where

$$\tilde{\mu}_{-i}(a_{-i} | b_i) := \sum_{b_{-i} \in B_{-i}} \psi_{-i}(b_{-i} | b_i) \prod_{j \in [n] \setminus i} \sigma_j(b_j)(a_j), \quad (5.3.2)$$

and $\text{supp}(\cdot)$ denotes the support of the distribution within the parentheses.

Definition 5.3.1. A randomized strategy profile σ is said to be a *mediated CPT Nash equilibrium* of a mediated game $\tilde{\Gamma} = (\Gamma, (B_i)_{i \in [n]})$ with respect to a mediator distribution $\psi \in \Delta(B)$ if

$$\sigma_i \in BR_i(\psi, \sigma) \text{ for all } i \in [n].$$

Let $\Sigma(\Gamma, (B_i)_{i \in [n]}, \psi)$ denote the set of all mediated CPT Nash equilibria of $\tilde{\Gamma} = (\Gamma, (B_i)_{i \in [n]})$ with respect to a mediator distribution $\psi \in \Delta(B)$.

For any mediator distribution $\psi \in \Delta(B)$, and any randomized strategy profile σ , let $\eta(\psi, \sigma) \in \Delta(A)$ be given by

$$\eta(\psi, \sigma)(a) := \sum_{b \in B} \psi(b) \prod_{i \in [n]} \sigma_i(b_i)(a_i). \quad (5.3.3)$$

Thus, $\eta(\psi, \sigma)$ gives the joint distribution over the action profiles of all the players corresponding to the randomized strategy σ and the mediator distribution ψ .

Definition 5.3.2. A probability distribution $\mu \in \Delta(A)$ is said to be a *mediated CPT correlated equilibrium* of a game Γ if there exist a signal system $(B_i)_{i \in [n]}$, a mediator distribution $\psi \in \Delta(B)$, and a mediated CPT Nash equilibrium $\sigma \in \Sigma(\Gamma, (B_i)_{i \in [n]}, \psi)$ such that $\mu = \eta(\psi, \sigma)$.

Consider an arbitrary mediated game $\tilde{\Gamma} = (\Gamma, (B_i)_{i \in [n]})$ with an arbitrary mediator distribution $\psi \in \Delta(B)$, where $B = \prod_{i=1}^n B_i$. If all the players choose to ignore the signals sent by the mediator, then the corresponding randomized strategy profile σ consists of constant functions $\sigma_i(b_i) \equiv \mu_i^*$. Further, as shown in Remark 5.A.1 in Appendix 5.A, it follows from Definition 3.2.4 and Definition 5.3.1 that the product probability distribution $\mu^* = \prod_{i \in [n]} \mu_i^*$ is a CPT Nash equilibrium of the game Γ iff σ is a mediated CPT Nash equilibrium of the mediated game $\tilde{\Gamma}$ with respect to the mediator distribution ψ . In particular, since every game Γ has at least one CPT Nash equilibrium, we see that every mediated game $\tilde{\Gamma}$ has at least one mediated CPT Nash equilibrium with respect to the mediator distribution ψ , for any mediator distribution ψ .

Let $D(\Gamma)$ denote the set of all mediated CPT correlated equilibria of a game Γ . By definition, $D(\Gamma)$ is the union over all signal systems $(B_i)_{i \in [n]}$ and mediator distributions $\psi \in \Delta(B)$ of $\{\eta(\psi, \sigma) : \sigma \in \Sigma(\Gamma, (B_i)_{i \in [n]}, \psi)\}$. When $B_i = A_i$ for all $i \in [n]$ and $\sigma = (\sigma_1, \dots, \sigma_n)$ is the deterministic strategy profile given, with an abuse of notation, by $\sigma_i(b_i)(a_i) = \mathbf{1}\{b_i = a_i\}$, one can check, see Remark 5.A.2 in Appendix 5.A, that $\sigma \in \Sigma(\Gamma, (A_i)_{i \in [n]}, \psi)$ iff $\psi \in C(\Gamma)$. In this case $\eta(\psi, \sigma) = \psi$ and so we have $C(\Gamma) \subset D(\Gamma)$. Under EUT, [5] proves that $D(\Gamma) = C(\Gamma)$. However, under CPT, this property, in general, does not hold true. Lemma 5.3.3 shows how $D(\Gamma)$ compares with $C(\Gamma)$.

For any $i, a_i, \tilde{a}_i \in A_i$, let $C(\Gamma, i, a_i, \tilde{a}_i)$ denote the set of all probability vectors $\pi_{-i} \in \Delta(A_{-i})$ such that

$$V_i(L_i(\pi_{-i}, a_i)) \geq V_i(L_i(\pi_{-i}, \tilde{a}_i)). \quad (5.3.4)$$

It is clear from the definition of CPT correlated equilibrium that, for a joint probability distribution $\mu \in C(\Gamma)$, provided $\mu_i(a_i) > 0$, the probability vector $\pi_{-i}(\cdot) = \mu_{-i}(\cdot | a_i) \in \Delta(A_{-i})$ should belong to $C(\Gamma, i, a_i, \tilde{a}_i)$ for all $\tilde{a}_i \in A_i$. Let

$$C(\Gamma, i, a_i) := \bigcap_{\tilde{a}_i \in A_i} C(\Gamma, i, a_i, \tilde{a}_i).$$

Notice that the set $C(\Gamma, i, a_i)$ is comprised of all probability vectors $\pi_{-i} \in \Delta(A_{-i})$ such that

$$V_i(L_i(\pi_{-i}, a_i)) \geq V_i(L_i(\pi_{-i}, \tilde{a}_i)), \forall \tilde{a}_i \in A_i.$$

In other words, $C(\Gamma, i, a_i)$ is the set of all probability distributions on the opponents' action profiles for which the lottery corresponding to action a_i gives the maximum CPT value to player i amongst all her actions. Now, for all i , define a subset $C(\Gamma, i) \subset \Delta(A)$, as follows:

$$C(\Gamma, i) := \{\mu \in \Delta(A) \mid \mu_{-i}(\cdot \mid a_i) \in C(\Gamma, i, a_i), \forall a_i \in \text{supp}(\mu_i)\}.$$

In words, $C(\Gamma, i)$ is the set of all probability distributions $\mu \in \Delta(A)$ such that, for every action a_i that has a positive probability under the marginal distribution μ_i , the conditional distribution $\mu_{-i}(\cdot \mid a_i)$ is such that the lottery corresponding to action a_i gives the maximum CPT value to player i amongst all her actions. Note that, since $V_i(L_i(\pi_{-i}, a_i))$ is a continuous function of π_{-i} , the sets $C(\Gamma, i, a_i, \tilde{a}_i)$, $C(\Gamma, i, a_i)$ and $C(\Gamma, i)$ are all closed.

Lemma 5.3.3. *For any game Γ , we have*

- (i) For all $i \in [n]$, $\overline{\text{co}}(C(\Gamma, i)) = \{\mu \in \Delta(A) \mid \mu_{-i}(\cdot \mid a_i) \in \overline{\text{co}}(C(\Gamma, i, a_i)), \forall a_i \in \text{supp}(\mu_i)\}$,
- (ii) $C(\Gamma) = \bigcap_{i \in [n]} C(\Gamma, i)$, and
- (iii) $D(\Gamma) = \bigcap_{i \in [n]} \overline{\text{co}}(C(\Gamma, i))$.

where $\overline{\text{co}}(S)$ denotes the convex hull of a set S .

We prove this in Appendix 5.D.

For the 2-person game Γ^* in Example 5.2.1, we observed that the set $C(\Gamma^*)$ is non-convex and hence $C(\Gamma^*) \neq D(\Gamma^*)$. If Γ is a 2×2 game, i.e., a game with 2 players, each having two actions, and both behaving according to CPT, then [108] prove that the sets $C(\Gamma, i)$, corresponding to both these players are convex, and hence also the set $C(\Gamma)$. From Lemma 5.3.3, we have the following result, having the flavor of the revelation principle:

Proposition 5.3.4. *If Γ is a 2×2 game, then the set of all CPT correlated equilibria is equal to the set of all mediated CPT correlated equilibria. \square*

In the context of mediated games, a strategy σ_i for player i is said to be *pure* if $\text{supp}(\sigma_i)$ is singleton and a strategy profile $\sigma = (\sigma_i)_{i \in [n]}$ is said to be a *pure strategy profile* if each σ_i is a pure strategy.

Remark 5.3.5. From the proof of Lemma 5.3.3, we observe that for any $\mu \in D(\Gamma)$, there exists a signal system $(B_i)_{i \in [n]}$ (of size $|B_i| = |A_i| \times |M_i| = |A|$), a mediator distribution $\psi \in \Delta(B)$, and a mediated CPT Nash equilibrium $\sigma \in \Sigma(\Gamma, (B_i)_{i \in [n]}, \psi)$ such that $\mu = \eta(\psi, \sigma)$ where σ is a pure strategy profile.

5.4 Calibrated Learning to Mediated CPT Correlated Equilibrium

Let ξ^t denote the empirical joint distribution of the action play up to step t . Formally,

$$\xi^t = \frac{1}{t} \sum_{\tau=1}^t e_{a^\tau},$$

where e_{a^τ} is an $|A|$ -dimensional vector with its a^τ -th component equal to 1 and the rest 0. We write the coordinates of ξ^t as $(\xi^t(a), a \in A)$. For each $i \in [n]$, we write $\xi_i^t := (\xi_i^t(a_i), a_i \in A_i)$ for the empirical distribution of the actions of player i . Thus ξ_i^t is the i -th marginal distribution corresponding to ξ^t . Similarly, for $i \in [n]$, $\xi_{-i}^t := (\xi_{-i}^t(a_{-i}|a_i), a \in A)$ are conditional distributions corresponding to ξ^t , where $\xi_{-i}^t(a_{-i}|a_i)$ is defined to be 0 when $\xi^t(a) = 0$.

Let the distance between a vector x and a set X in the same Euclidean space be given by

$$d(x, X) = \inf_{x' \in X} \|x - x'\|,$$

where $\|x\|$ denotes the standard Euclidean norm of x . We say that a sequence $(x^t, t \geq 1)$ converges to a set X if the following holds:

$$\lim_{t \rightarrow \infty} d(x^t, X) = 0.$$

Theorem 5.4.1. *Assume that the assessment schemes and best reaction maps of the players are such that if each player at each step plays the best reaction to her assessment then each player is calibrated with respect to the sequence of action profiles of the other players. Then the empirical joint distribution of action play ξ^t converges to the set of mediated CPT correlated equilibria.*

We prove this theorem in Appendix 5.E.

Remark 5.4.2. In the proof of Theorem 5.4.1, in fact, we prove the following stronger statement: If player i 's assessments are calibrated with respect to the sequence of action profiles of her opponents and she chooses the best reaction to her assessments at every step, then the joint empirical distribution of action play converges to the set $\overline{\text{co}}(C(\Gamma, i))$.

Now the question remains whether each player i can make assessments that are guaranteed to be calibrated no matter what strategies her opponents use. But this has nothing to do whether the players have EUT or CPT preferences, and has been answered in the affirmative by Theorem 3 in [50]. To be precise, at each step t , the player i predicts a distribution $\mu_{-i}^t \in \Delta(A_{-i})$ by drawing one from a distribution over the space of distributions $\Delta(A_{-i})$, determined by the history H^{t-1} (which we recall is given by the sequence of action profiles of all the players over the steps up to $t-1$) and a random seed U_i^t , where the seeds $(U_i^t, t \geq 1)$ are i.i.d. and independent of the randomizations, if any, used by the other players. The rule

by which this probability distribution is created as a function of H^{t-1} and U_i^t is assumed to be common knowledge to all the players. The assessment of player i at step t is then the realization of this random choice. Lumping together the opponents of player i as Nature from the point of view of this player, at each step t , Nature can be assumed to have access not only to the history H^{t-1} but also to the realizations of the past seed values $(U_i^1, \dots, U_i^{t-1})$, so Nature knows the assessments of the player i from steps 1 to $t-1$. Crucially, while Nature now knows the distribution of the assessment of player i at time t , Nature does not know the realization of this assessment till the next time step. In this scenario (referred to as the *adaptive adversary* scenario in [49]), a strategy for Nature is comprised of Nature playing an action at step t by drawing one randomly from a distribution on her set of actions (i.e. the set A_{-i} of action profiles of the opponents of player i) based on the information available to her at this step, namely H^{t-1} and $(U_i^1, \dots, U_i^{t-1})$. The calibrated learning result proved in [49] says that there exists such a randomized forecasting scheme on the part of player i such that, no matter what randomized strategy Nature employs as above, we have

$$\sum_{q \in Q^t} |\rho(q, y, t) - q(y)| \frac{N(q, t)}{t} \rightarrow 0, \text{ as } t \rightarrow \infty, \quad (5.4.1)$$

for all $y \in A_{-i}$, almost surely (over the random seeds of player i and the randomization in Nature's strategy).¹⁷ Here, as in equation (5.2.1), Q^t denotes the set of probability distributions in $\Delta(A_{-i})$ actually predicted by player i up to step t .

Combining this result with Theorem 5.4.1 we have the following corollary, proved in Appendix 5.F.

Corollary 5.4.3. *There exist a randomized assessment scheme and a best reaction map for each player such that, if each player predicts her assessments according to her scheme and plays the best reaction to her assessments, then it is almost surely true (over the randomization in the randomized assessment schemes for the players) that each player is calibrated with respect to the sequence of action profiles of her opponents, and hence the empirical distribution of action play converges to the set of mediated CPT correlated equilibria.*

We now show that, in a certain sense, the set $D(\Gamma)$ is the smallest possible extension of the set $C(\Gamma)$ that guarantees convergence of the empirical distribution of action play to this set, when all the players have assessment schemes and best reaction maps such that when each player plays the best reaction to her assessment at each step the player is calibrated with respect to the sequence of action profiles of her opponents. In particular, we claim the following.

Proposition 5.4.4. *For all games Γ such that the sets $C(\Gamma, i, a_i)$, $i \in [n]$, $a_i \in A_i$ do not have any isolated points, if $\mu \in D(\Gamma)$, then there exists an assessment scheme and a best reaction map for each player such that if each player plays her best reaction to her assessment at each step then each player's assessments are calibrated with respect to the sequence of action profiles of her opponents and the empirical distribution of action play converges to μ .*

See Appendix 5.G for a proof of this proposition. The following proposition (proved in Appendix 5.H) shows under some technical conditions on the value function of each player that, for generic games Γ , the sets $C(\Gamma, i, a_i)$, $i \in [n]$, $a_i \in A_i$, do not have any isolated points. For any player i , we know that the value function $v_i^{r_i}(x)$ is a strictly increasing continuous function. Let the open interval $Y_i \subset \mathbb{R}$ denote the range of $v_i^{r_i}$, and let λ_i^* denote the push forward measure of the Lebesgue measure on \mathbb{R} with respect to the function $v_i^{r_i}$. Let $\hat{\lambda}_i$ denote the Lebesgue measure restricted to the interval Y_i . We will require that the function $v_i^{r_i}$ is such that $\lambda_i^* \ll \hat{\lambda}_i$ (i.e., the measure λ_i^* is absolutely continuous with respect to the measure $\hat{\lambda}_i$). Since the function $v_i^{r_i}$ is strictly increasing, its inverse function $(v_i^{r_i})^{-1} : Y_i \rightarrow \mathbb{R}$ is well defined. We have $\lambda_i^* \ll \hat{\lambda}_i$ if and only if the function $(v_i^{r_i})^{-1}$ is absolutely continuous.

Proposition 5.4.5. *For any fixed CPT features $r_i, v_i^{r_i}, w_i^\pm$ such that $(v_i^{r_i})^{-1}$ is absolutely continuous, and a fixed action set A_i for each of the players $i \in [n]$ (here, we assume $n > 1$ and $|A_i| > 1, \forall i \in [n]$), the set of all games Γ for which there exists a player $i \in [n]$ and an action $a_i \in A_i$ such that the set $C(\Gamma, i, a_i)$ has an isolated point is a null set with respect to the Lebesgue measure λ on the space of payoffs $(x_i(a), a \in A, i \in [n])$, viewed as an $n \times |A|$ -dimensional Euclidean space.*

5.5 No-Regret Learning and CPT Correlated Equilibrium

The randomized forecasting scheme proposed in [49] generates a probability distribution on the space of assessments of player i . Player i draws her assessment from this distribution and then plays her best reaction. This two step process gives rise to a randomized strategy for player i at each step. Together with Remark 5.4.2 we get that, no matter what strategies the opponents play, player i can guarantee that the empirical distribution of action play converges almost surely to the set $\overline{\text{co}}(C(\Gamma, i))$.

Under EUT, player i has a strategy that guarantees the almost sure convergence of the empirical distribution of action play to the set $C(\Gamma, i)$. This convergence is related to the notion of no-regret learning. We now describe this approach. Suppose that, at step t , player i imagines replacing action a_i by action \tilde{a}_i , every time she played action a_i in the past. Assuming the actions of the other players did not change, her payoff would become $x_i(\tilde{a}_i, a_{-i}^\tau)$ for all $\tau \leq t$ such that $a_i^\tau = a_i$, instead of $x_i(a_i, a_{-i}^\tau)$, while for all $\tau \leq t$ such that $a_i^\tau \neq a_i$ it will continue to be $x_i(a_i^\tau)$. We define the resulting *CPT regret* of player i for having played action a_i instead of action \tilde{a}_i as

$$K_i^t(a_i, \tilde{a}_i) := \xi_i^t(a_i) \mathcal{R}_i \left[\left\{ \left(\xi_{-i}^t(a_{-i} | a_i), x_i(\tilde{a}_i, a_{-i}), x_i(a_i, a_{-i}) \right) \right\}_{a_{-i} \in A_{-i}} \right], \quad (5.5.1)$$

where

$$\mathcal{R}_i [\{(\nu_l, \hat{z}_l, z_l)\}_{l=1}^m] := V_i(\{(\nu_l, \hat{z}_l)\}_{l=1}^m) - V_i(\{(\nu_l, z_l)\}_{l=1}^m) \quad (5.5.2)$$

is the difference in the CPT values of the lotteries $\{(\nu_l, \hat{z}_l)\}_{l=1}^m$ and $\{(\nu_l, z_l)\}_{l=1}^m$. We associate player i with CPT regrets $\{K_i^t(a_i, \tilde{a}_i), a_i, \tilde{a}_i \in A_i, a_i \neq \tilde{a}_i\}$ at each step t . Under EUT, this simplifies to

$$K_i^t(a_i, \tilde{a}_i) = \frac{1}{t} \sum_{\tau \leq t: a_i^\tau = a_i} [x_i(\tilde{a}_i, a_{-i}^\tau) - x_i(a_i^\tau)], \quad (5.5.3)$$

in agreement with the definition given in [60].

The following proposition (proved in Appendix 5.I) shows the connection between regrets and correlated equilibrium.

Proposition 5.5.1. *Let $(a^t)_{t \geq 1}$ be a sequence of action profiles played by the players. Then $\limsup_{t \rightarrow \infty} K_i^t(a_i, \tilde{a}_i) \leq 0$, for every $i \in [n]$ and every $a_i, \tilde{a}_i \in A_i, a_i \neq \tilde{a}_i$, if and only if the sequence of empirical distributions ξ^t converges to the set $C(\Gamma)$ of CPT correlated equilibrium.*

Player i is said to have a no-regret learning strategy if, irrespective of the strategies of the other players, her regrets satisfy

$$P \left(\limsup_{t \rightarrow \infty} K_i^t(a_i, \tilde{a}_i) \leq 0 \right) = 1, \text{ for every } a_i, \tilde{a}_i \in A_i, a_i \neq \tilde{a}_i.$$

This is equivalent to asking if the vector of regrets $(K_i^t(a_i, \tilde{a}_i), a_i, \tilde{a}_i \in A_i, a_i \neq \tilde{a}_i)$, converges to the nonpositive orthant almost surely. This is related to the concept of approachability, the setup for which is as follows. Consider a repeated two player game, where now at step t , if the row player and the column player play actions \hat{a}_{row}^t and \hat{a}_{col}^t respectively, then the row player receives a vector payoff $\vec{x}(\hat{a}_{row}^t, \hat{a}_{col}^t)$ instead of a scalar payoff. A subset S is said to be approachable by the row player if she has a (randomized) strategy such that, no matter how the column player plays, we have

$$\lim_{t \rightarrow \infty} d \left(\frac{1}{t} \sum_{\tau=1}^t \vec{x}(\hat{a}_{row}^\tau, \hat{a}_{col}^\tau), S \right) = 0, \text{ almost surely.}$$

Blackwell's approachability theorem [16] establishes that a convex closed set S is approachable if and only if every halfspace \mathcal{H} containing S is approachable.

[60] cast the repeated game with stage game Γ in the above setup as a two player repeated game where player i is the row player and the opponents together form the column player. Let $\vec{x}(\hat{a}_i, \hat{a}_{-i})$ be the vector payoff when player i plays action \hat{a}_i and the others play \hat{a}_{-i} , with components given by

$$\vec{x}_{a_i, \tilde{a}_i}(\hat{a}_i, \hat{a}_{-i}) = \begin{cases} x_i(\tilde{a}_i, \hat{a}_{-i}) - x_i(a_i, \hat{a}_{-i}) & \text{if } a_i = \hat{a}_i, \\ 0 & \text{otherwise,} \end{cases}$$

for all $a_i, \tilde{a}_i \in A_i, a_i \neq \tilde{a}_i$. Under EUT, the average vector payoff of the row player corresponds to the regret of player i (see equation (5.5.3)). Hart and Mas-Colell [60] show that

the nonpositive orthant is approachable for the row player and hence player i has a no-regret learning strategy. Under CPT, if the average vector payoffs were to match the regrets of player i , then the vector payoffs for the row player at step t would need to depend on the empirical distribution of action play up to step t . Indeed, the component corresponding to the pair (a_i, \tilde{a}_i) of the vector payoff for the row player at step t when player i plays action \hat{a}_i and the others play \hat{a}_{-i} would need to match the difference

$$(t+1)K_i^{t+1}(a_i, \tilde{a}_i) - tK_i^t(a_i, \tilde{a}_i).$$

This difference depends on the empirical distribution of action play up to step t , and hence in general changes with t . This suggests that there might be difficulties in adapting the approach in [60] to study no-regret learning strategies under CPT.

The following example shows that, under CPT, approachability of the nonpositive orthant need not hold. In other words, it can happen under CPT that at least one of the players does not have a no-regret learning strategy.

Example 5.5.2. Consider the 2-player repeated game from Example 5.2.1. Recall the following distributions on player 2's actions: $\sigma_{odd} = (0.5, 0, 0.5, 0)$, $\sigma_{even} = (0, 0.5, 0, 0.5)$ and $\sigma_{unif} = (0.25, 0.25, 0.25, 0.25)$. We observed that player 1's action 1 is not a best response to σ_{odd} and σ_{even} and player 1's action 0 is not a best response to σ_{unif} . For an integer $T > 2$, consider the following strategy for player 2:

- play mixed strategy σ_{odd} at step 1,
- play mixed strategy σ_{even} at step 2,
- play mixed strategy σ_{odd} at steps $2T^k < t \leq T^{k+1}$, for $k \geq 0$,
- play mixed strategy σ_{even} at steps $T^{k+1} < t \leq 2T^{k+1}$, for $k \geq 0$.

The rest of this section will be devoted to proving that player 1 cannot have a no-regret learning strategy. In particular, we will prove the following in Appendix 5.K:

Proposition 5.5.3. *In the above example, for a suitable choice of $T, \tilde{\delta} > 0$ and $\tilde{\epsilon} > 0$, there exists an integer k_0 such that no matter what learning strategy player 1 uses, for all $k \geq k_0$ we have*

$$P(\bar{K}^k > \tilde{\epsilon}) > \tilde{\delta},$$

where

$$\bar{K}^k := [K_1^{T^{k+1}}(1, 0)]^+ + [K_1^{2T^{k+1}}(0, 1)]^+ + [K_1^{2T^{k+1}}(1, 0)]^+, \quad (5.5.4)$$

using the notation $[\cdot]^+ := \max\{\cdot, 0\}$. Here, for actions a_i and \tilde{a}_i of player 1, $K_1^t(a_i, \tilde{a}_i)$ are the CPT regrets of player 1 at step t , as defined in equation (5.5.1).

Consider the subsequence of steps $(t_{odd}^l)_{l \geq 1}$ when player 2 played σ_{odd} . Let $\nu_{odd}^l(a_1, a_2)$ denote the empirical distribution over those times of the action profile (a_1, a_2) , where $a_1 \in \{0, 1\}$, $a_2 \in \{I, III\}$, i.e.

$$\nu_{odd}^l(a_1, a_2) := \frac{1}{l} \sum_{u=1}^l \mathbf{1}\{a^{t_{odd}^u} = (a_1, a_2)\}. \quad (5.5.5)$$

Similarly, consider the sequence of steps $(t_{even}^l)_{l \geq 1}$ when player 2 played σ_{even} . Let $\nu_{even}^l(a_1, a_2)$ denote the empirical distribution over those times of the action profile (a_1, a_2) , where $a_1 \in \{0, 1\}$, $a_2 \in \{II, IV\}$, i.e.

$$\nu_{even}^l(a_1, a_2) := \frac{1}{l} \sum_{u=1}^l \mathbf{1}\{a^{t_{even}^u} = (a_1, a_2)\}. \quad (5.5.6)$$

Lemma 5.5.4. For any $\delta > 0$, there exists an integer $l_\delta > 1$, such that for all $l \geq l_\delta$, we have

$$P(|\nu_{odd}^l(0, I) - \nu_{odd}^l(0, III)| < \delta) > 1 - \delta, \quad (5.5.7)$$

$$P(|\nu_{odd}^l(1, I) - \nu_{odd}^l(1, III)| < \delta) > 1 - \delta, \quad (5.5.8)$$

$$P(|\nu_{even}^l(0, II) - \nu_{even}^l(0, IV)| < \delta) > 1 - \delta, \quad (5.5.9)$$

$$P(|\nu_{even}^l(1, II) - \nu_{even}^l(1, IV)| < \delta) > 1 - \delta. \quad (5.5.10)$$

The proof of Lemma 5.5.4 can be found in Appendix 5.J.

For a vector $q \in \mathbb{R}^S$ and $\epsilon > 0$, let $[q]_\epsilon := \{\tilde{q} \in \mathbb{R}^S : |\tilde{q}(s) - q(s)| < \epsilon, \forall s \in S\}$ denote the set of all vectors strictly within ϵ of q in the sup norm. Select positive constants $\epsilon_3, c_3, \epsilon_2, c_2, \epsilon_1, c_1$ as follows:

- Let $\epsilon_3 < 1$ and c_3 be such that for the indicated regret we have

$$\mathcal{R}_1[\{(\mu(\cdot), x_1(1, \cdot), x_1(0, \cdot))\}] > c_3, \quad (5.5.11)$$

for all probability distributions $\mu \in [\sigma_{unif}]_{\epsilon_3}$ (such constants exist because action 0 is not a best response to σ_{unif}). Let

$$\delta_3 := \epsilon_3/2. \quad (5.5.12)$$

Note that $\delta_3 < 1/2$.

- Let $\epsilon_2 < 1$ and c_2 be such that for the indicated regret we have

$$\mathcal{R}_1[\{(\mu(\cdot), x_1(0, \cdot), x_1(1, \cdot))\}] > c_2, \quad (5.5.13)$$

for all probability distributions $\mu \in [\sigma_{even}]_{\epsilon_2}$ (such constants exist because action 1 is not a best response to σ_{even}). Let

$$\delta_2 := \epsilon_2 \delta_3 / 4. \quad (5.5.14)$$

Note that $\delta_2 < 1/8$.

- Let $\epsilon_1 < 0.5$ and c_1 be such that for the indicated regret we have

$$\mathcal{R}_1 [\{(\mu(\cdot), x_1(0, \cdot), x_1(1, \cdot))\}] > c_1, \quad (5.5.15)$$

for all probability distributions $\mu \in [\sigma_{odd}]_{\epsilon_1}$ (such constants exist because action 1 is not a best response to σ_{odd}). Let

$$\delta_1 := \epsilon_1 \delta_2. \quad (5.5.16)$$

Note that $\delta_1 < 1/16$.

Let $T > 2/\delta_1$ and k_0 be such that

$$T^{k_0+1} > \max \left\{ t_{odd}^{l_{\delta_1}}, t_{odd}^{l_{\delta_1}}, t_{even}^{l_{\delta_1}}, t_{even}^{l_{\delta_1}} \right\}, \quad (5.5.17)$$

where l_{δ_1} is such that the inequalities in Lemma 5.5.4 hold for $\delta = \delta_1$.

For $k \geq 0$, let f_1^{k+1} denote the fraction of times player 2 plays σ_{even} up to step $t = T^{k+1}$. From the definition of the strategy of player 2, we have

$$f_1^{k+1} < \frac{2T^k}{T^{k+1}} = \frac{2}{T}. \quad (5.5.18)$$

Similarly, for $k \geq 0$, let f_2^{k+1} denote the fraction of times player 2 plays σ_{even} up to step $t = 2T^{k+1}$. We have

$$f_2^{k+1} = \frac{T^{k+1} + \frac{T^{k+1}-1}{T-1}}{2T^{k+1}} \in \left[\frac{1}{2}, \frac{1}{2} + \frac{1}{T} \right], \quad (5.5.19)$$

where the last inclusion follows from the assumption that $T > 2$. Note that

$$f_2^{k+1} = 1/2 + f_1^{k+1}/2.$$

Next, for $k \geq 0$, let

$$f_3^{k+1} := \xi_1^{T^{k+1}}(0), \quad (5.5.20)$$

i.e. the fraction of times player 1 plays action 0 up to step $t = T^{k+1}$, and let

$$f_4^{k+1} := 2\xi_1^{2T^{k+1}}(0) - \xi_1^{T^{k+1}}(0), \quad (5.5.21)$$

i.e. the fraction of times player 1 plays action 0 among the steps from $T^{k+1} + 1$ to $2T^{k+1}$. Note that f_3^{k+1} and f_4^{k+1} are random variables, in contrast with f_1^{k+1} and f_2^{k+1} .

We will establish the proof of Proposition 5.5.3 in stages through several lemmas. In the next couple of paragraphs we first outline our proof strategy.

Depending on the strategy of player 1, we have two possibilities, either $P(f_3^{k+1} < 1 - \delta_2) > 1/4$ or $P(f_3^{k+1} < 1 - \delta_2) \leq 1/4$. In the former case, in Lemma 5.K.3, we show that the empirical distribution $\xi^{T^{k+1}}(1, \cdot)$ is restricted to be of a certain type with significant probability, conditioned on $\{f_3^{k+1} < 1 - \delta_2\}$. The purpose of this lemma is to show that the conditional distribution $\xi_{-1}^{T^{k+1}}(\cdot|1)$ is close to σ_{odd} . We explain this in Lemma 5.K.4, and use

it to establish that player 1 has a significant regret at step T^{k+1} for not having played action 0 whenever she played action 1 up to that step, i.e. $K_1^{T^{k+1}}(1,0)$ is considerable.

In the latter case, in Lemma 5.K.5, we show that the distribution $\xi^{2T^{k+1}}$ is restricted to be of a certain type with significant probability, conditioned on $\{f_3^{k+1} \geq 1 - \delta_2\}$. We then consider two cases depending on f_4^{k+1} , which was defined in equation (5.5.21). If f_4^{k+1} is less than $1 - \delta_3$, then, in Lemma 5.K.6, we show that the conditional distribution $\xi_{-1}^{2T^{k+1}}(\cdot|1)$ is similar to σ_{even} and hence player 1 suffers from a significant regret at step $2T^{k+1}$ for not having played action 0 whenever she played action 1 up to that step, i.e. $K_1^{2T^{k+1}}(1,0)$ is considerable. If f_4^{k+1} is greater than or equal to $1 - \delta_3$, then, in Lemma 5.K.7, we show that the conditional distribution $\xi_{-1}^{2T^{k+1}}(\cdot|0)$ is similar to σ_{unif} and hence player 1 suffers from a significant regret at step $2T^{k+1}$ for not having played action 1 whenever she played action 0 up to that step, i.e. $K_1^{2T^{k+1}}(0,1)$ is considerable. Finally, we can combine these results to show that player 1 faces some regret either at step T^{k+1} or $2T^{k+1}$ for all $k \geq k_0$, and hence the regret vector of player 1 never converges to the nonpositive orthant.

5.6 Summary

We studied how some of the results from the theory of learning in games are affected when the players in the game have cumulative prospect theoretic preferences. For example, we saw that the notion of mediated CPT correlated equilibrium arising from mediated games is more appropriate than the notion of CPT correlated equilibrium while studying the convergence of the empirical distribution of action play, in particular for calibrated learning schemes. One can ask similar questions with respect to other learning schemes such as *follow the perturbed leader* [53], *fictitious play* [22], etc. We leave this for future work. In general, it seems that the results from the theory of learning in games continue to hold under CPT with slight modifications.

We also observed that the revelation principle does not hold under CPT. In the next chapter, we will see the implications of this to mechanism design.

Appendix

5.A Notions of equilibrium

In this appendix, we explore the relationship between the different notions of equilibrium for a finite n -person normal form game Γ with CPT players, organizing our observations into a sequence of remarks. For convenience, we first briefly recall the four notions of equilibrium that played a role in the discussion in this chapter. A CPT correlated equilibrium of the game Γ , see Definition 3.2.3, is an element of $\Delta(A)$. A CPT Nash equilibrium of the game Γ , see Definition 3.2.4, is an element of $\Delta^*(A)$. Given a signal system $(B_i)_{i \in [n]}$ and a mediator distribution $\psi \in \Delta(B)$, where $B := \prod_{i=1}^n B_i$, a mediated CPT Nash equilibrium of the

mediated game $\tilde{\Gamma} := (\Gamma, (B_i)_{i \in [n]})$ with respect to the mediator distribution ψ , see Definition 5.3.1, is a randomized strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$, where $\sigma_i : B_i \rightarrow \Delta(A_i)$. A mediated CPT correlated equilibrium of the game Γ , see Definition 5.3.2, is an element of $\Delta(A)$.

Remark 5.A.1. Let $\mu := \prod_{i=1}^n \mu_i \in \Delta^*(A)$ be a CPT Nash equilibrium of the game Γ . Then, for every signal system $(B_i)_{i \in [n]}$ and mediator distribution $\psi \in \Delta(B)$, the randomized strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$, where $\sigma_i : B_i \rightarrow \Delta(A_i)$ is the constant function given by $\sigma_i(b_i) = \mu_i$ for all $b_i \in B_i$, is a mediated CPT Nash equilibrium of the mediated game $\tilde{\Gamma} := (\Gamma, (B_i)_{i \in [n]})$ with respect to the mediator distribution ψ . Conversely, if σ is defined in terms of $\mu \in \Delta^*(A)$ as above and σ is a mediated CPT Nash equilibrium of the mediated game $\tilde{\Gamma} := (\Gamma, (B_i)_{i \in [n]})$ with respect to the mediator distribution ψ , then μ is a CPT Nash equilibrium of the game Γ .

To see this, note that for the strategy profile σ , for all $b_i \in B_i$, we have $\tilde{\mu}_{-i}(a_{-i}|b_i) = \prod_{j \neq i} \mu_j(a_j)$ for all $a_{-i} \in A_{-i}$, where $\tilde{\mu}_{-i}(a_{-i}|b_i)$ is as defined in equation (5.3.2). Hence $\sigma_i \in BR_i(\psi, \sigma)$, where $BR_i(\psi, \sigma)$ is as defined in equation (5.3.1), iff $\mu_i \in BR_i(\mu)$, where $BR_i(\mu)$ is as defined in equation (3.2.4). This establishes the claim.

Remark 5.A.2. Every CPT correlated equilibrium of the game Γ is a mediated CPT correlated equilibrium of the game Γ . Namely $C(\Gamma) \subset D(\Gamma)$.

To see this, let $\mu \in C(\Gamma)$. Consider the signal system $(A_i)_{i \in [n]}$ (i.e. take $B_i = A_i$ for all $i \in [n]$) with the mediator distribution μ , and consider the deterministic strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ given, with an abuse of notation, by $\sigma_i(b_i) = \mathbf{1}\{b_i = a_i\}$. Note that $\eta(\psi, \sigma)$, as defined in equation (5.3.3), equals μ . Since $\mu \in C(\Gamma)$, it verifies the condition in equation (3.2.3), which then implies that $\sigma_i \in BR_i(\psi, \sigma)$, where $\psi = \mu$ and $BR_i(\psi, \sigma)$ is as defined in equation (5.3.1). This implies that $\mu \in D(\Gamma)$.

Remark 5.A.3. Suppose the mediator distribution ψ is of product form, which we write as $\psi \in \Delta^*(B)$. Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a mediated CPT Nash equilibrium of the mediated game $\tilde{\Gamma} := (\Gamma, (B_i)_{i \in [n]})$ with respect to the mediator distribution ψ . Let $\mu := \eta(\psi, \sigma)$, as defined in equation (5.3.3). Note that we will have $\mu \in \Delta^*(A)$. A simple calculation shows that $\tilde{\mu}_i(a_{-i}|b_i) = \prod_{j \neq i} \mu_j(a_j)$ for all $i \in [n]$, $b_i \in B_i$, and $a_{-i} \in A_{-i}$, where $\tilde{\mu}_i(a_{-i}|b_i)$ is as defined in equation (5.3.2). Thus $\sigma_i \in BR_i(\psi, \sigma)$ iff for all $b_i \in \text{supp}(\psi_i)$ we have $\sigma_i(b_i) \in BR_i(\mu)$. This, in turn, is equivalent to $\mu_i \in BR_i(\mu)$. This characterizes the mediated CPT Nash equilibria of a mediated game $\tilde{\Gamma} := (\Gamma, (B_i)_{i \in [n]})$ with respect to product form mediator distributions $\psi \in \Delta^*(B)$ in terms the CPT Nash equilibria of the game Γ .

Remark 5.A.4. In [101], the authors showed that for any finite n-person game the Nash equilibria all lie on the boundary of the set of correlated equilibria. We extended this result to the CPT setting and showed that all the CPT Nash equilibria lie on the boundary of the set of CPT correlated equilibria in Chapter 3. It is natural to ask whether the CPT Nash equilibria in fact lie on the boundary of the set of all mediated CPT correlated equilibria. We know this is true for any 2×2 game Γ , since $C(\Gamma) = D(\Gamma)$ for such games. However, it is not known if this property holds in general for all finite n -person CPT games, and we leave this for future work.

5.B Beyond Fixed Reference Points

CPT differs from EUT in several ways:

- (i) The reference point r divides the outcomes into two domains – gains ($x > r$) and losses ($x < r$) – and governs the shape of the value function v^r . Recall that $v^r(r) = 0$, and typically $v^r(x)$ is convex in the losses frame and concave in the gains frame, and it is steeper in the losses frame than in the gains frame.
- (ii) The probability weighting functions w^\pm govern the probabilistic sensitivity of an individual for gains and losses, respectively.
- (iii) The cumulative functional form of the CPT value function based on the reference point, the value function and the probability weighting functions governs the CPT preferences.

Throughout this chapter we have assumed that the reference points r_i are exogenous and hence fixed, for all players i . In Example 5.2.1 and Example 5.5.2, the reference points of the players were set to be zero, all the outcomes were non-negative (hence, in the gains domain), and the value functions were set to be the identity functions. Thus, the results from these two examples are purely an effect of the probability weighting functions and the cumulative functional form of the CPT values (equation (1.3.9)). The two results, namely:

1. the lack of convergence of the empirical distribution of action play to the set of CPT correlated equilibria when each player's assessments are calibrated with respect to the sequence of action profiles of the other players and she evaluates her best reaction in accordance to CPT preferences, and
2. the existence of instances of repeated games where a player does not have a no-regret learning strategy under CPT preferences,

carry over to other restrictive versions of CPT such as rank-dependent utility theory (RDU) [114] and Yaari's dual theory [135].

Although our framework with exogenous reference points is rich enough to capture loss aversion, diminishing sensitivity to returns, and the cumulative functional form of the CPT value function that depends on the reference point and treats gains and losses separately, it is interesting to consider reference points that are determined endogenously. For instance, in the definition of CPT Nash equilibrium and CPT correlated equilibrium given in Chapter 3, we allow the reference point r_i to depend on the equilibrium distribution $\mu \in \Delta(A)$ given by a function $\hat{r}_i : \Delta(A) \rightarrow \mathbb{R}$, for each player i . (For other works with endogenous reference points, see [121] and [78].)

In general, there is no consensus in the literature on how the players update their reference points at each step of a repeated game (see, for example, [13, 82, 81, 77, 57]). Taking the viewpoint that the reference point indicates the expectations of the player for the decision-making problem at hand, we assume that the reference point of player i at step t is determined

by her assessment $\mu_{-i}^t \in \Delta(A_{-i})$ over the actions of the other players. Let this be given by a function

$$\tilde{r}_i : \Delta(A_{-i}) \rightarrow \mathbb{R},$$

for each player i (we assume that \tilde{r}_i is same across all the steps). Let $V_i^{r_i}(L_i)$ denote the CPT value evaluated by player i for lottery L_i when her reference point is r_i . At step t , if player i 's assessment of the actions of the other players is μ_{-i}^t , then she evaluates the lottery corresponding to each of her actions using the CPT value function $V_i^{\tilde{r}_i(\mu_{-i}^t)}$ based on the reference point $\tilde{r}_i(\mu_{-i}^t)$ and the value function $v_i^{\tilde{r}_i(\mu_{-i}^t)}$. As before, we assume that she prefers to play an action that gives her the maximum CPT value (now incorporating the updated reference point $\tilde{r}_i(\mu_{-i}^t)$). Accordingly, for any assessment $\mu_{-i} \in \Delta(A_{-i})$, we assume that player i has a corresponding action that gives her the maximum CPT value (with any arbitrary but fixed tie breaking rule if multiple actions correspond to the maximum CPT value). We call this her best reaction for her assessment μ_{-i} .

We now consider a repeated game where each player's assessments are calibrated with respect to the sequence of action profiles of the other players and she evaluates her best reaction in accordance with CPT preferences corresponding to the reference point determined by her assessment at each step. To study the convergence of the empirical distribution of play, we need to modify the definitions of CPT correlated equilibrium and mediated CPT correlated equilibrium to account for the reference points based on functions \tilde{r}_i , for all players i .

A joint probability distribution $\mu \in \Delta(A)$ will be called a CPT correlated equilibrium (with reference points \tilde{r}_i) of a game Γ if it satisfies the following inequalities for all i , and for all $a_i, \tilde{a}_i \in A_i$, such that $\mu_i(a_i) > 0$:

$$V_i^{\tilde{r}_i(\mu_{-i}(\cdot|a_i))}(L_i(\mu_{-i}(a_{-i}|a_i), a_i)) \geq V_i^{\tilde{r}_i(\mu_{-i}(\cdot|\tilde{a}_i))}(L_i(\mu_{-i}(a_{-i}|\tilde{a}_i), \tilde{a}_i)).$$

Let $\tilde{C}(\Gamma)$ denote the set of all CPT correlated equilibria (with reference points given by $\tilde{r}_i : \Delta(A_{-i}) \rightarrow \mathbb{R}$). Notice that this definition of CPT correlated equilibrium is different from the definition given in Chapter 3, where the reference points are determined by the functions $\hat{r}_i : \Delta(A) \rightarrow \mathbb{R}$. These two definitions coincide when the reference points are fixed, i.e. \tilde{r}_i and \hat{r}_i are constant functions.

We can also define the notion of CPT Nash equilibrium (with reference points \tilde{r}_i) as follows: For any mixed strategy profile $\mu \in \Delta^*(A)$, where $\Delta^*(A)$ is defined in equation (3.2.2), the reference point for each player i is given by $\tilde{r}_i(\mu_{-i})$, where $\mu_{-i} = \prod_{j \neq i} \mu_j$. If player i decides to deviate and play a mixed strategy $\tilde{\mu}_i$ while the rest of the players continue to play μ_{-i} , then define the average CPT value for player i by

$$\tilde{\mathcal{A}}_i(\tilde{\mu}_i, \mu_{-i}) := \sum_{a_i \in A_i} \tilde{\mu}_i(a_i) V_i^{\tilde{r}_i(\mu_{-i})}(L_i(\mu_{-i}, a_i)).$$

The best response set of player i to a mixed strategy profile $\mu \in \Delta^*(A)$ can now be defined as

$$\begin{aligned} \widetilde{BR}_i(\mu) &:= \left\{ \mu_i^* \in \Delta(A_i) \mid \forall \tilde{\mu}_i \in \Delta(A_i), \mathcal{A}_i(\mu_i^*, \mu_{-i}) \geq \mathcal{A}_i(\tilde{\mu}_i, \mu_{-i}) \right\} \\ &= \left\{ \mu_i^* \in \Delta(A_i) \mid \text{supp}(\mu_i^*) \subset \arg \max_{a_i \in A_i} V_i^{\tilde{r}_i(\mu_{-i})}(L_i(\mu_{-i}, a_i)) \right\}. \end{aligned}$$

A mixed strategy profile $\mu^* \in \Delta^*(A)$ is a CPT Nash equilibrium (with reference points \tilde{r}_i) of Γ iff

$$\mu_i^* \in \widetilde{BR}_i(\mu^*) \text{ for all } i.$$

We observe that CPT Nash equilibrium (with reference points \tilde{r}_i) as defined here is a special case of CPT Nash equilibrium as defined in [74]. To see this, notice that, for any mixed strategy profile $\mu \in \Delta^*(A)$, the reference point \hat{r}_i is a function of the mixed strategies $\mu_i, \forall i \in [n]$, whereas \tilde{r}_i is a function of the mixed strategies of the opponents, namely, $\mu_j, \forall j \neq i$. Thus, when we restrict our attention to mixed strategy profiles $\mu \in \Delta^*(A)$, the reference points \tilde{r}_i are a special case of the reference points \hat{r}_i . From this observation and the existence of CPT Nash equilibrium (with reference points \hat{r}_i) established in [Theorem 1]keskin2016equilibrium, we conclude that, for any game Γ , there exists a CPT Nash equilibrium (with reference points \tilde{r}_i).

Further, it can be shown that the set of all CPT Nash equilibria of a game Γ with CPT players with reference points given by $\tilde{r}_i : \Delta(A_{-i}) \rightarrow \mathbb{R}$ is equal to $\tilde{C}(\Gamma) \cap \Delta^*(A)$. The proof of this is identical, with the obvious modifications, to the one given in Proposition 1 of [74] for the notion of CPT correlated equilibrium considered there, i.e. where the reference points are determined by the functions $\hat{r}_i : \Delta(A) \rightarrow \mathbb{R}$. Thus, we get that the set $\tilde{C}(\Gamma)$ is non-empty.

In a mediated game, the best response set of player i as defined in Equation 3.2.4, can be modified as follows: For a mediated game $\tilde{\Gamma} = (\Gamma, (B_i)_{i \in [n]})$, the best response set of player i (with reference points \tilde{r}_i) to a randomized strategy profile σ and a mediator distribution ψ is defined as

$$\begin{aligned} \widetilde{BR}_i(\psi, \sigma) &:= \left\{ \sigma_i^* : B_i \rightarrow \Delta(A_i) \mid \text{for all } b_i \in \text{supp}(\psi_i), \right. \\ &\quad \left. \text{supp}(\sigma_i^*(b_i)) \subset \arg \max_{a_i \in A_i} V_i^{\tilde{r}_i(\tilde{\mu}_{-i}(a_{-i}|b_i))} \left(\{ \tilde{\mu}_{-i}(a_{-i}|b_i), x_i(a_i, a_{-i}) \}_{a_{-i} \in A_{-i}} \right) \right\}, \end{aligned}$$

where $\tilde{\mu}_{-i}(a_{-i}|b_i)$ is as defined in equation (5.3.2), namely,

$$\tilde{\mu}_{-i}(a_{-i}|b_i) = \sum_{b_{-i} \in B_{-i}} \psi_{-i}(b_{-i}|b_i) \prod_{j \in [n] \setminus i} \sigma_j(b_j)(a_j),$$

A randomized strategy profile σ is said to be a mediated CPT Nash equilibrium (with reference points \tilde{r}_i) of a mediated game $\tilde{\Gamma} = (\Gamma, (B_i)_{i \in [n]})$ with respect to a mediator distribution $\psi \in \Delta(B)$ if

$$\sigma_i \in \widetilde{BR}_i(\psi, \sigma) \text{ for all } i \in [n].$$

A probability distribution $\mu \in \Delta(A)$ is said to be a mediated CPT correlated equilibrium (with reference points \tilde{r}_i) of a game Γ if there exist a signal system $(B_i)_{i \in [n]}$, a mediator distribution $\psi \in \Delta(B)$, and a mediated CPT Nash equilibrium (with endogenous reference points \tilde{r}_i) σ such that $\mu = \eta(\psi, \sigma)$. Let $\tilde{D}(\Gamma)$ denote the set of all such equilibria. It can be seen that $\tilde{D}(\Gamma)$ contains $\tilde{C}(\Gamma)$, by considering the special case where the signal sets B_i are equal to the respective action sets A_i , the mediator distribution is in $\tilde{C}(\Gamma)$, and the strategy of each player is to play the action suggested by the mediator. Hence $\tilde{D}(\Gamma)$ is nonempty.

As before, we can define for any $i, a_i, \tilde{a}_i \in A_i$, the set $\tilde{C}(\Gamma, i, a_i, \tilde{a}_i)$ of all probability vectors $\pi_{-i} \in \Delta(A_{-i})$ such that

$$V_i^{\tilde{r}_i(\pi_{-i})}(L_i(\pi_{-i}, a_i)) \geq V_i^{\tilde{r}_i(\pi_{-i})}(L_i(\pi_{-i}, \tilde{a}_i)).$$

Let

$$\tilde{C}(\Gamma, i, a_i) := \bigcap_{\tilde{a}_i \in A_i} \tilde{C}(\Gamma, i, a_i, \tilde{a}_i),$$

and

$$\tilde{C}(\Gamma, i) := \{\mu \in \Delta(A) \mid \mu_{-i}(\cdot \mid a_i) \in \tilde{C}(\Gamma, i, a_i), \forall a_i \in \text{supp}(\mu_i)\}.$$

Similarly to Lemma 5.3.3, we can show that

$$\tilde{C}(\Gamma) = \bigcap_{i \in [n]} \tilde{C}(\Gamma, i),$$

and

$$\tilde{D}(\Gamma) = \bigcap_{i \in [n]} \overline{\text{co}}(\tilde{C}(\Gamma, i)).$$

Lemma 5.3.3 and its proof extends verbatim if we replace $C(\Gamma, i, a_i)$, $C(\Gamma, i)$, and $D(\Gamma)$ by $\tilde{C}(\Gamma, i, a_i)$, $\tilde{C}(\Gamma, i)$, and $\tilde{D}(\Gamma)$, respectively.

Similarly to Theorem 5.4.1, we can show that in a repeated game when each player's assessments are calibrated with respect to the sequence of action profiles of the other players and at each step she plays the best reaction to her assessment (the best reaction map now depends on the function \tilde{r}_i), then the empirical distribution of action play converges to the set $\tilde{D}(\Gamma)$. The proof of Theorem 5.4.1 extends verbatim to this setting if we replace $C(\Gamma, i, a_i)$, $C(\Gamma, i)$, and $D(\Gamma)$ by $\tilde{C}(\Gamma, i, a_i)$, $\tilde{C}(\Gamma, i)$, and $\tilde{D}(\Gamma)$, respectively, and interpret the set $R_i(a_i) \subset \Delta(A_{-i})$ of all joint distributions μ_{-i} for which action a_i is player i 's best reaction as per the CPT preferences of player i with endogenous reference point \tilde{r}_i .

Along similar lines, we can extend Proposition 5.4.4 and its proof by replacing $C(\Gamma, i, a_i)$, $C(\Gamma, i)$, and $D(\Gamma)$ by $\tilde{C}(\Gamma, i, a_i)$, $\tilde{C}(\Gamma, i)$, and $\tilde{D}(\Gamma)$, respectively. The analog of Proposition 5.4.5, however, remains aloof, and we leave it for future work.

5.C Generalized Signal Spaces

We now allow the signal set B_i to be an arbitrary Polish space (a complete separable metric space) for all $i \in [n]$. The product spaces $B := \prod_{i \in [n]} B_i$ and $B_{-i} := \prod_{j \neq i} B_j$, for all $i \in [n]$,

are then also Polish spaces because a countable product of Polish spaces is a Polish space. Let $\mathcal{B}_i, \mathcal{B}$ and \mathcal{B}_{-i} denote the σ -algebra of Borel sets on the spaces B_i, B and B_{-i} respectively. Let the mediator be characterized by a probability distribution ψ on (B, \mathcal{B}) . Let ψ_i denote the marginal probability distribution on B_i induced by ψ . Let $\psi_{-i} : B_i \times \mathcal{B}_{-i} \rightarrow [0, 1]$ be a function which satisfies:

1. $\psi_{-i}(b_i, \cdot)$ is a probability distribution on $(B_{-i}, \mathcal{B}_{-i})$, for all $b_i \in B_i$,
2. $\psi_{-i}(\cdot, X)$ is a measurable function on (B_i, \mathcal{B}_i) , for all $X \in \mathcal{B}_{-i}$,
3. for all $X \in \mathcal{B}_{-i}$ and $Y \in \mathcal{B}_i$,

$$\psi(Y \times X) = \int_Y \psi_{-i}(y, X) \psi_i(dy). \quad (5.C.1)$$

The function ψ_{-i} is called a *regular conditional probability*. For a proof of its existence, see [29, Theorem 1] (this theorem needs to be used in the framework of [29, Example 2]).

Let a randomized strategy for any player i be given by a measurable function $\sigma_i : B_i \rightarrow \Delta(A_i)$ with respect to the Borel σ -algebra on $\Delta(A_i)$, and let $\sigma = (\sigma_1, \dots, \sigma_n)$ denote the randomized strategy profile as before. Let $\sigma_{-i} := \prod_{j \neq i} \sigma_j : B_{-i} \rightarrow \Delta(A_{-i})$. Let $\nu_{-i}(b_i)$ be the push forward probability distribution of $\psi_{-i}(b_i, \cdot)$ with respect to the function σ_{-i} , and let

$$\tilde{\mu}_{-i}(a_{-i}|b_i) := \int_{\Delta(A_{-i})} p(a_{-i}) \nu_{-i}(b_i)(dp). \quad (5.C.2)$$

Note that $\tilde{\mu}_{-i}(\cdot|b_i) \in \Delta(A_{-i})$. Let $\nu(\psi, \sigma)$ be the push forward probability distribution of ψ with respect to the function $\sigma := \prod_{i \in [n]} \sigma_i : B \rightarrow \Delta(A)$, and let

$$\eta(\psi, \sigma)(a) := \int_{\Delta(A)} p(a) \nu(\psi, \sigma)(dp). \quad (5.C.3)$$

Note that $\eta(\psi, \sigma) \in \Delta(A)$.

Let the best response set of player i to a randomized strategy profile σ and a mediator distribution ψ be given by

$$BR_i(\psi, \sigma) := \left\{ \sigma_i^* : B_i \rightarrow \Delta(A_i) \text{ a measurable function} \mid \text{for all } b_i \in \text{supp}(\psi_i), \right. \\ \left. \text{supp}(\sigma_i^*(b_i)) \subset \arg \max_{a_i \in A_i} V_i \left(\{ \tilde{\mu}_{-i}(a_{-i}|b_i), x_i(a_i, a_{-i}) \}_{a_{-i} \in A_{-i}} \right) \right\}, \quad (5.C.4)$$

where $\text{supp}(\psi_i)$ is the smallest closed set $Y \subset B_i$ with $\psi_i(B_i \setminus Y) = 0$.

We can now define, exactly as in Definition 5.3.1, the notion of a mediated CPT Nash equilibrium for the mediated game $\tilde{\Gamma} := (\Gamma, (B_i)_{i \in [n]})$ with respect to a probability distribution ψ on (B, \mathcal{B}) , where now $(B_i, \mathcal{B}_i)_{i \in [n]}$ are arbitrary Polish spaces. Let $\Sigma^*(\Gamma, (B_i)_{i \in [n]}, \psi)$ denote the set of such mediated CPT Nash equilibria. We can also define, exactly as in

Definition 5.3.2, the notion of a mediated CPT correlated equilibrium (which is a probability distribution in $\Delta(A)$, as before) in this extended setting where the signal spaces are allowed to be arbitrary Polish spaces. Let $D^*(\Gamma)$ denote the set of mediated CPT correlated equilibria in this extended sense. Let $C(\Gamma, i, a_i)$ and $C(\Gamma, i)$ be defined as before.

Lemma 5.C.1. *For any game Γ , we have*

$$D^*(\Gamma) \subset \bigcap_{i \in [n]} \overline{\text{co}}(C(\Gamma, i)).$$

Proof. Let $\mu \in D^*(\Gamma)$. Then there exists a signal system comprised of Polish spaces $(B_i, \mathcal{B}_i)_{i \in [n]}$, a mediator distribution ψ which is a probability distribution on (B, \mathcal{B}) , and a mediated CPT Nash equilibrium $\sigma \in \Sigma^*(\Gamma, (B_i)_{i \in [n]}, \psi)$ such that $\mu = \eta(\psi, \sigma)$. Fix $i \in [n]$. For $b_i \in \text{supp}(\psi_i)$ and $a_i \in \text{supp}(\sigma_i(b_i))$, we have $\tilde{\mu}_{-i}(\cdot | b_i) \in C(\Gamma, i, a_i)$, from equations (5.C.4) and (5.3.4). Let a_i be such that $\mu_i(a_i) > 0$. We have

$$\mu_{-i}(\cdot | a_i) = \int_{B_i} \frac{\sigma_i(b_i)(a_i)}{\mu_i(a_i)} \tilde{\mu}_{-i}(\cdot | b_i) \psi_i(db_i).$$

Also, since σ is the product function $\prod_{i \in [n]} \sigma_i$ and μ is the push forward probability distribution of ψ with respect to σ , we have that μ_i is the push forward probability distribution of ψ_i with respect to the function σ_i , i.e.

$$\mu_i(a_i) = \int_{B_i} \sigma_i(b_i)(a_i) \psi_i(db_i).$$

Since the set $\overline{\text{co}}(C(\Gamma, i, a_i))$ is closed, we have $\mu_{-i}(\cdot | a_i) \in \overline{\text{co}}(C(\Gamma, i, a_i))$. Since this holds for all $i \in [n]$, we have $\mu = \eta(\psi, \sigma) \in \bigcap_{i \in [n]} \overline{\text{co}}(C(\Gamma, i))$. This completes the proof. \square

Since a finite set B_i is a Polish space with respect to the discrete topology, we have $D(\Gamma) \subset D^*(\Gamma)$. From the above lemma and Lemma 5.3.3 we have $D^*(\Gamma) = D(\Gamma)$. Hence, it is enough to restrict our attention to signals B_i that are finite sets. In fact, it suffices to restrict attention to signal sets B_i of size at most $|A|$ (see Remark 5.3.5).

5.D Proof of Lemma 5.3.3

Fix $i \in [n]$. Note that, since the sets $C(\Gamma, i)$ and $C(\Gamma, i, a_i)$ for each $a_i \in A_i$ are closed, the convex hulls of these sets are closed. Suppose $\mu = \lambda\mu^1 + (1 - \lambda)\mu^2$ where $\mu^1, \mu^2 \in C(\Gamma, i)$ and $0 < \lambda < 1$. If $a_i \in \text{supp}(\mu_i)$, then one of the following three cases holds:

Case 1 [$a_i \in \text{supp}(\mu_i^1)$, $a_i \in \text{supp}(\mu_i^2)$]. Then, $\mu_{-i}^1(\cdot | a_i), \mu_{-i}^2(\cdot | a_i) \in C(\Gamma, i, a_i)$ and we have,

$$\mu_{-i}(\cdot | a_i) = \frac{\lambda\mu_i^1(a_i)\mu_{-i}^1(\cdot | a_i) + (1 - \lambda)\mu_i^2(a_i)\mu_{-i}^2(\cdot | a_i)}{\lambda\mu_i^1(a_i) + (1 - \lambda)\mu_i^2(a_i)}.$$

Let $\theta = (\lambda\mu_i^1(a_i))/(\lambda\mu_i^1(a_i) + (1 - \lambda)\mu_i^2(a_i))$. Then $\mu_{-i}(\cdot | a_i) = \theta\mu_{-i}^1(\cdot | a_i) + (1 - \theta)\mu_{-i}^2(\cdot | a_i)$ and hence $\mu_{-i}(\cdot | a_i) \in \overline{\text{co}}(C(\Gamma, i, a_i))$.

Case 2 [$a_i \in \text{supp}(\mu_i^1)$, $a_i \notin \text{supp}(\mu_i^2)$] Here $\mu_{-i}(\cdot|a_i) = \mu_{-i}^1(\cdot|a_1)$. Hence $\mu_{-i}(\cdot|a_i) \in C(\Gamma, i, a_i)$.

Case 3 [$a_i \notin \text{supp}(\mu_i^1)$, $a_i \in \text{supp}(\mu_i^2)$] This can be handled similarly to case 2.

Also, the above argument can be easily extended to when μ is a convex combination of any finite number of distributions. Since all our sets are compact subsets of finite dimensional Euclidean spaces, Caratheodory's theorem applies, and hence we need to consider only finite convex combinations.

This shows that the set on the left hand side is contained in the set on the right hand side of the equation in (i) for the given fixed $i \in [n]$.

To prove the inclusion in the other direction, fix $i \in [n]$ and let μ belong to the set on the right hand side of the equation in (i). If $a_i \in \text{supp}(\mu_i)$, then $\mu_{-i}(\cdot|a_i)$ is a linear combination of $|A_{-i}|$ joint distributions (allowing repetitions), call them

$$\zeta_{-i,a_i}^1, \dots, \zeta_{-i,a_i}^{m_i}, \dots, \zeta_{-i,a_i}^{|A_{-i}|} \in C(\Gamma, i, a_i),$$

with coefficients $\theta_{i,a_i}^{m_i}$, $1 \leq m_i \leq |A_{-i}|$ respectively (where $0 < \theta_{i,a_i}^{m_i} \leq 1$ for all $1 \leq m_i \leq |A_{-i}|$ can be ensured because we allow repetitions). For each $\zeta_{-i,a_i}^{m_i}$, construct a distribution $\mu_{i,a_i}^{m_i} \in \Delta(A)$ by $\mu_{i,a_i}^{m_i}(\tilde{a}_i, \tilde{a}_{-i}) = 1\{\tilde{a}_i = a_i\}\zeta_{-i,a_i}^{m_i}(\tilde{a}_{-i})$. Then $\mu_{i,a_i}^{m_i} \in C(\Gamma, i)$. Let $\lambda_{i,a_i}^{m_i} := \mu_i(a_i)\theta_{i,a_i}^{m_i}$, for all i, m_i, a_i such that $\mu_i(a_i) > 0$. One can now check that $\mu = \sum_{m_i, a_i} \lambda_{i,a_i}^{m_i} \mu_{i,a_i}^{m_i}$ for the given fixed $i \in [n]$. Thus $\mu \in \overline{\text{co}}(C(\Gamma, i))$.

Statement (ii) follows directly from the definition of CPT correlated equilibrium.

For statement (iii), let $\mu \in \Delta(A)$ be such that $\mu \in \overline{\text{co}}(C(\Gamma, i))$ for all i . For any a_i such that $\mu_i(a_i) > 0$, by (i), we have, $\mu_{-i}(\cdot|a_i) \in \overline{\text{co}}(C(\Gamma, i, a_i))$. As above, let $\mu_{-i}(\cdot|a_i)$ be a convex combination of $|A_{-i}|$ joint distributions (allowing repetitions), call them

$$\zeta_{-i,a_i}^1, \dots, \zeta_{-i,a_i}^{m_i}, \dots, \zeta_{-i,a_i}^{|A_{-i}|} \in C(\Gamma, i, a_i),$$

with coefficients $\theta_{i,a_i}^{m_i}$, $1 \leq m_i \leq |A_{-i}|$ respectively (where $0 < \theta_{i,a_i}^{m_i} \leq 1$ for all $1 \leq m_i \leq |A_{-i}|$ can be ensured because we allow repetitions). For all i , let $B_i := A_i \times M_i$, where $M_i := \{1, \dots, |A_{-i}|\}$. Let the mediator distribution be given by

$$\psi((a_1, m_1), \dots, (a_n, m_n)) = \begin{cases} \frac{\mu(a) \prod_{i=1}^n \{\theta_{i,a_i}^{m_i} \zeta_{-i,a_i}^{m_i}(a_{-i})\}}{\sum_{\tilde{m}_i, i \in [n]} \prod_{i=1}^n \{\theta_{i,a_i}^{\tilde{m}_i} \zeta_{-i,a_i}^{\tilde{m}_i}(a_{-i})\}}, & \text{if } \mu(a) > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (5.D.1)$$

It is useful to note that

$$\sum_{\tilde{m}_i, i \in [n]} \prod_{i=1}^n \{\theta_{i,a_i}^{\tilde{m}_i} \zeta_{-i,a_i}^{\tilde{m}_i}(a_{-i})\} = \prod_{i=1}^n \mu_{-i}(a_{-i}|a_i), \quad (5.D.2)$$

and that, for every $i \in [n]$,

$$\psi_i((a_i, m_i)) := \sum_{(a_j, m_j), j \in [n] \setminus i} \psi((a_1, m_1), \dots, (a_n, m_n)) = \mu_i(a_i) \theta_{i,a_i}^{m_i}. \quad (5.D.3)$$

Let the strategy for each player i be

$$\sigma_i(a_i, m_i)(\tilde{a}_i) = \begin{cases} 1, & \text{if } \tilde{a}_i = a_i, \\ 0, & \text{otherwise.} \end{cases} \quad (5.D.4)$$

From equations (5.3.3), (5.D.1) and (5.D.4) we have

$$\begin{aligned} \eta(\psi, \sigma)(a) &= \sum_{(\tilde{a}_i, m_i) \in B_i, i \in [n]} \psi((\tilde{a}_1, m_1), \dots, (\tilde{a}_n, m_n)) \prod_{i \in [n]} \sigma_i((\tilde{a}_i, m_i))(a_i) \\ &= \sum_{m_i, i \in [n]} \psi((a_1, m_1), \dots, (a_n, m_n)) \\ &= \mu(a) \sum_{m_i, i \in [n]} \frac{\prod_{i=1}^n \{\theta_{i, a_i}^{m_i} \zeta_{-i, a_i}^{m_i}(a_{-i})\}}{\sum_{\tilde{m}_i, i \in [n]} \prod_{i=1}^n \{\theta_{i, a_i}^{\tilde{m}_i} \zeta_{-i, a_i}^{\tilde{m}_i}(a_{-i})\}} \\ &= \mu(a). \end{aligned}$$

From equations (5.3.2), (5.D.1), (5.D.2), (5.D.3) and (5.D.4) we have

$$\begin{aligned} \tilde{\mu}_{-i}(a_{-i} | (a_i, m_i)) &= \sum_{(\tilde{a}_j, m_j) \in B_j, j \in [n] \setminus i} \psi_{-i}((\tilde{a}_j, m_j), j \in [n] \setminus i) | (a_i, m_i) \\ &\quad \times \prod_{j \in [n] \setminus i} \sigma_j((\tilde{a}_j, m_j))(a_j) \\ &= \sum_{m_j, j \in [n] \setminus i} \psi_{-i}(((a_j, m_j), j \in [n] \setminus i) | (a_i, m_i)) \\ &= \frac{\sum_{m_j, j \in [n] \setminus i} \psi((a_1, m_1), \dots, (a_n, m_n))}{\psi_i((a_i, m_i))} \\ &= \zeta_{-i, a_i}^{m_i}(a_{-i}). \end{aligned}$$

Thus we have $\tilde{\mu}_{-i}(\cdot | (a_i, m_i)) \in C(\Gamma, i, a_i)$. Hence $\mu \in D(\Gamma)$. We have established that $\bigcap_{i \in N} \overline{\text{co}}(C(\Gamma, i)) \subset D(\Gamma)$.

For the other direction of statement (iii), let $\mu \in D(\Gamma)$. Then there exists a signal system $(B_i)_{i \in [n]}$, a mediator distribution $\psi \in \Delta(B)$, and a mediated CPT Nash equilibrium $\sigma \in \Sigma(\Gamma, (B_i)_{i \in [n]}, \psi)$ such that $\mu = \eta(\psi, \sigma)$. Fix $i \in [n]$. For $b_i \in \text{supp}(\psi_i)$ and $a_i \in \text{supp}(\sigma_i(b_i))$, we have $\tilde{\mu}_{-i}(\cdot | b_i) \in C(\Gamma, i, a_i)$, from equations (5.3.1) and (5.3.4). But

$$\mu_{-i}(\cdot | a_i) = \sum_{b_i \in \text{supp}(\psi_i)} \frac{\psi_i(b_i) \sigma_i(b_i)(a_i)}{\mu_i(a_i)} \tilde{\mu}_{-i}(\cdot | b_i).$$

Hence $\mu_{-i}(\cdot | a_i) \in \overline{\text{co}}(C(\Gamma, i, a_i))$. Since this holds for all $i \in [n]$, we have $\mu = \eta(\psi, \sigma) \in \bigcap_{i \in [n]} \overline{\text{co}}(C(\Gamma, i))$. This completes the proof. \square

5.E Proof of Theorem 5.4.1

Consider the sequence of empirical distributions ξ^t . Since the simplex $\Delta(A)$ of all joint distributions over action profiles is a compact set, every such sequence has a convergent subsequence. Thus, it is enough to show that the limit of any convergent subsequence of ξ^t is in $D(\Gamma)$. Let ξ^{t_k} be such a convergent subsequence and denote its limit by $\hat{\xi}$.

Let a_i be an action of player i such that $\hat{\xi}_i(a_i) > 0$. Let $R_i(a_i) \subset \Delta(A_{-i})$ be the set of all joint distributions μ_{-i} for which action a_i is player i 's best reaction. Note that $R_i(a_i)$ forms a partition of the simplex $\Delta(A_{-i})$. Let $\mu_{-i}^t \in \Delta(A_{-i})$ denote player i 's assessment at step t , and let Q_i^t denote the set of assessments made by her up to step t . Since $\hat{\xi}_i(a_i) > 0$, there exists an integer $k_0 \geq 1$ and an $\epsilon > 0$ such that, for all $k \geq k_0$, we have $\xi_i^{t_k}(a_i) > \epsilon$. For all $k \geq k_0$, we have

$$\begin{aligned}
\xi_{-i}^{t_k}(a_{-i}|a_i)\xi_i^{t_k}(a_i)t_k &= \sum_{\substack{\tau \leq t_k \\ \text{s.t. } \mu_{-i}^\tau \in R_i(a_i)}} \mathbf{1}\{a_{-i}^\tau = a_{-i}\} \\
&= \sum_{q \in R_i(a_i) \cap Q_i^{t_k}} \sum_{\substack{\tau \leq t_k \\ \text{s.t. } \mu_{-i}^\tau = q}} \mathbf{1}\{a_{-i}^\tau = a_{-i}\} \\
&= \sum_{q \in R_i(a_i) \cap Q_i^{t_k}} \rho(q, a_{-i}, t_k) N(q, t_k) \\
&= \sum_{q \in R_i(a_i) \cap Q_i^{t_k}} q(a_{-i}) N(q, t_k) \\
&+ \sum_{q \in R_i(a_i) \cap Q_i^{t_k}} (\rho(q, a_{-i}, t_k) - q(a_{-i})) N(q, t_k).
\end{aligned}$$

Using

$$\xi_i^{t_k}(a_i)t_k = \sum_{q \in R_i(a_i) \cap Q_i^{t_k}} N(q, t_k),$$

we get, for all $k \geq k_0$,

$$\begin{aligned}
\xi_{-i}^{t_k}(a_{-i}|a_i) &= \frac{\sum_{q \in R_i(a_i) \cap Q_i^{t_k}} q(a_{-i}) N(q, t_k)}{\sum_{q \in R_i(a_i) \cap Q_i^{t_k}} N(q, t_k)} \\
&+ \frac{1}{\xi_i^{t_k}(a_i)} \sum_{q \in R_i(a_i) \cap Q_i^{t_k}} (\rho(q, a_{-i}, t_k) - q(a_{-i})) \frac{N(q, t_k)}{t_k}.
\end{aligned}$$

Since player i is calibrated with respect to the sequence of action profiles of her opponents, the second term in the last expression goes to zero as $k \rightarrow \infty$ (here, we use the fact that

$\xi_i^{t_k}(a_i) > \epsilon > 0$ for all $k \geq k_0$). Further, we have, for all $k \geq 1$,

$$\frac{\sum_{q \in R_i(a_i) \cap Q_i^{t_k}} q N(q, t_k)}{\sum_{q \in R_i(a_i) \cap Q_i^{t_k}} N(q, t_k)} \in \overline{\text{co}}(R_i(a_i)).$$

Taking the limit as $k \rightarrow \infty$ we have, $\hat{\xi}_{-i}(\cdot | a_i) \in \overline{\text{co}}(R_i(a_i))$, where $\overline{\text{co}}(\cdot)$ denotes the closed convex hull. Note that $R_i(a_i) \subset C(\Gamma, i, a_i)$ and $C(\Gamma, i, a_i)$ is closed. Thus $\hat{\xi}_{-i}(\cdot | a_i) \in \overline{\text{co}}(C(\Gamma, i, a_i))$ for all $a_i \in A_i$ such that $\hat{\xi}_i(a_i) > 0$. By Lemma 5.3.3, we have $\hat{\xi} \in \overline{\text{co}}(C(\Gamma, i))$, and since this is true for all players i , we have $\hat{\xi} \in D(\Gamma)$. \square

5.F Proof of Corollary 5.4.3

Let player i be the forecaster and let all the opponents together form Nature from the point of view of the player. Thus Nature's action set is A_{-i} . By the [49] result, there exists a randomized assessment scheme for player i such that, whatever the randomized strategy that Nature uses, the sequence of assessments of player i is calibrated almost surely with respect to the sequence of actions of Nature. Let player i use such a randomized scheme to determine her assessments. From Remark 5.4.2, it follows that the empirical distribution of play converges to the set $\overline{\text{co}}(C(\Gamma, i))$ almost surely. If each player follows such a strategy, then we get almost sure convergence to $D(\Gamma)$. \square

5.G Proof of Proposition 5.4.4

Since $\mu \in D(\Gamma)$, as noted in Remark 5.3.5, there exists a signal system $(B_i)_{i \in [n]}$ where B_i can be identified with $A_i \times A_{-i}$, a mediator distribution $\psi \in \Delta(B)$, and a mediated CPT Nash equilibrium $\sigma \in \Sigma(\Gamma, (B_i)_{i \in [n]}, \psi)$ such that $\mu = \eta(\psi, \sigma)$, where σ is a pure strategy profile. With an abuse of notation, let $\sigma_i(b_i)$ denote the unique element in the support of $\sigma_i(b_i)$. Let (b^1, b^2, \dots) be a sequence of signal profiles such that the empirical distribution of these signal profiles converges to ψ and such that $\psi(b^t) > 0$ for all $t \geq 1$. At step t , let player i predict her assessment $\tilde{\mu}_{-i}(\cdot | b_i)$ (as defined in equation (5.3.2)) and play $\sigma_i(b_i)$. The sequence of assessments of each player is calibrated with respect to the sequence of action profiles of her opponents. To see this, fix a player i , let $q \in \Delta(A_{-i})$ be one of the assessments made by her, and let $G = \{b_i \in B_i | \tilde{\mu}_{-i}(\cdot | b_i) = q\}$. Let $t^k(b_i)$ denote the step when player i receives signal b_i for the k th time. By construction, the empirical average of the action profiles of the opponents of player i over the steps $(t^k(b_i))_{k \geq 1}$ converges to $\tilde{\mu}_{-i}(\cdot | b_i)$. As a result, the empirical average of the action profiles of the opponents of player i over the steps when player i receives a signal $b_i \in G$ converges to q . Since this holds for any assessment q made by player i , her assessments are calibrated. Further, by construction, the empirical distribution of action play converges to μ .

If $\tilde{\mu}_{-i}(\cdot | b_i) = \tilde{\mu}_{-i}(\cdot | \tilde{b}_i)$ implies $\sigma_i(b_i) = \sigma_i(\tilde{b}_i)$, for all $b_i, \tilde{b}_i \in B_i, i \in [n]$, then we can define $\sigma_i(b_i)$ as the best reaction to the assessment $\tilde{\mu}_{-i}(\cdot | b_i)$ and the claim is proved. \square

there exist b_i, \tilde{b}_i such that $\tilde{\mu}_{-i}(\cdot|b_i) = \tilde{\mu}_{-i}(\cdot|\tilde{b}_i)$ but $\sigma_i(b_i) \neq \sigma_i(\tilde{b}_i)$, then there is a problem in defining the best reaction to the assessment $\tilde{\mu}_{-i}(\cdot|b_i)$. We now describe a way to get around such a situation, analogous to the scheme used in [50] to resolve the same kind of issue. Let $\zeta_{-i}^* := \tilde{\mu}_{-i}(\cdot|b_i) = \tilde{\mu}_{-i}(\cdot|\tilde{b}_i)$ and let $a_i^* := \sigma_i(b_i) \neq \sigma_i(\tilde{b}_i)$. By the assumption that the set $C(\Gamma, i, a_i^*)$ does not have any isolated points, there exists a sequence $(\hat{\zeta}_{-i}^l)_{l \geq 1}$ of distinct probability distributions in $C(\Gamma, i, a_i^*)$ such that $\hat{\zeta}_{-i}^l \rightarrow \zeta_{-i}^*$ and $(\hat{\zeta}_{-i}^l)_{l \geq 1}$ are all distinct from the distributions $(\tilde{\mu}_{-i}(\cdot|b_i), \forall b_i \in B_i)$. Further, let the sequence $(\hat{\zeta}_{-i}^l)_{l \geq 1}$ be such that $|\hat{\zeta}_{-i}^l(a_{-i}) - \zeta_{-i}^*(a_{-i})| < 1/l$, for all $a_{-i} \in A_{-i}$, i.e. $\hat{\zeta}_{-i}^l$ is within $1/l$ of ζ_{-i}^* in the sup norm, for all $l \geq 1$. We will now replace the assessments ζ_{-i}^* at the steps $(t^k(b_i))_{k \geq 1}$ by the assessments $(\hat{\zeta}_{-i}^l)_{l \geq 1}$, with each $\hat{\zeta}_{-i}^l$ repeated sufficiently many times that the empirical distribution of the action profiles of the opponents over the steps that player i 's assessment is $\hat{\zeta}_{-i}^l$ is within $1/l$ of ζ_{-i}^* in the sup norm. To achieve this, start by replacing the assessment at step $t^1(b_i)$ by $\hat{\zeta}_{-i}^1$. Next replace the assessments at steps $t^k(b_i), k = 2, 3, \dots$ with $\hat{\zeta}_{-i}^2$ until the empirical distribution of the action profiles of the opponents over these steps is within $1/2$ of ζ_{-i}^* in the sup norm. In general, keep replacing the assessments at steps $t^k(b_i)$ with $\hat{\zeta}_{-i}^l$ until the empirical distribution of the action profiles of the opponents over these steps is within $1/l$ of ζ_{-i}^* in the sup norm, and then switch to replacing by $\hat{\zeta}_{-i}^{l+1}$. Note that each assessment $\hat{\zeta}_{-i}^l$ will be used only for a finite number of steps since the empirical distribution of the action profiles of the opponents over the steps $(t^k(b_i))_{k \geq 1}$ converges to ζ_{-i}^* . Thus, the empirical distribution of the action profiles of the opponents over the steps when player i makes assessment $\hat{\zeta}_{-i}^l$ is within $2/l$ of $\hat{\zeta}_{-i}^l$ in the sup norm. We know that if a sequence of probability distributions $(s_t)_{t \geq 1}$ on A_{-i} converges to a probability distribution s on A_{-i} , then the sequence of the running averages $S_t = (1/t) \sum_{\tau=1}^t s_\tau, t \geq 1$, also converges to s . Using this fact, we observe that the sequence of player i 's assessments continues to be calibrated with respect to the sequence of action profiles of her opponents even after the above replacement. Since the assessments $\{\hat{\zeta}_{-i}^l\}$ are distinct from the assessments $(\tilde{\mu}_{-i}(\cdot|b_i), \forall b_i \in B_i)$, we can define action a_i^* as the best reaction to $\hat{\zeta}_{-i}^l$ for all $l \geq 1$. The above trick can be used to resolve all instances where $\tilde{\mu}_{-i}(\cdot|b_i) = \tilde{\mu}_{-i}(\cdot|\tilde{b}_i)$ but $\sigma_i(b_i) \neq \sigma_i(\tilde{b}_i)$. Each time taking the corresponding sequence $\{\hat{\zeta}_{-i}^l\}$ distinct from all previously used assessments. This solves the problem of defining the best reaction map of each player and completes the proof. \square

5.H Proof of Proposition 5.4.5

For each of the players $i \in [n]$, let us fix the CPT features $r_i, v_i^{r_i}, w_i^\pm$ such that $(v_i^{r_i})^{-1}$ is absolutely continuous. We also fix the action set A_i for each of the players $i \in [n]$. Since n and $|A_i|, \forall i$ are finite, it is enough to show that for any fixed $i \in [n]$ and $a_i \in A_i$ the set of all games Γ for which the set $C(\Gamma, i, a_i)$ has an isolated point is a null set. Since the set of all games for which any two payoffs of player i are equal, i.e. $x_i(a) = x_i(\tilde{a}), a \neq \tilde{a}$, is a null set, we can restrict our attention to games where all the payoffs for player i corresponding to her playing a_i are distinct. Let $(\pi_i(1), \pi_i(2), \dots, \pi_i(|A_{-i}|))$ be a permutation of A_{-i} such

that

$$x_i(a_i, \pi_i(1)) > x_i(a_i, \pi_i(2)) > \cdots > x_i(a_i, \pi_i(|A_{-i}|)).$$

Suppose we fix $x_j(a) \in \mathbb{R}$ for all $j \neq i$, and $x_i(\tilde{a}_i, a_{-i}) \in \mathbb{R}$ for all $\tilde{a}_i \neq a_i, a_{-i} \in A_{-i}$. Then the game Γ is completely determined by the vector of payoffs $(x_i(a_i, a_{-i}))_{a_{-i} \in A_{-i}}$. Let S denote the set of all $(x_i(a_i, a_{-i}))_{a_{-i} \in A_{-i}}$ for which the set $C(\Gamma, i, a_i)$ has isolated points. We will show that S is a null set with respect to the Lebesgue measure on $\mathbb{R}^{|A_{-i}|}$. Then, by Tonelli's theorem, we have the required result.

Recall that $Y_i \subset \mathbb{R}$ denotes the range of $v_i^{r_i}$ and that Y_i is an open interval because $v_i^{r_i}$ is assumed to be continuous and strictly increasing on \mathbb{R} . Also recall that λ_i^* is the measure on Y_i that is the push forward of the Lebesgue measure on \mathbb{R} under $v_i^{r_i}$, $\hat{\lambda}_i$ denotes Lebesgue measure restricted to Y_i , and that the assumption that $(v_i^{r_i})^{-1}$ is absolutely continuous implies that λ_i^* is absolutely continuous with respect to $\hat{\lambda}_i$. Consider the function $f : \mathbb{R}^{|A_{-i}|} \rightarrow Y_i^{|A_{-i}|}$ given by

$$f((x_i(a_i, a_{-i}))_{a_{-i} \in A_{-i}}) := (v_i^{r_i}(x_i(a_i, a_{-i})))_{a_{-i} \in A_{-i}}$$

Let $y_i(a_{-i}) := v_i^{r_i}(x_i(a_i, a_{-i})) \in Y_i$ for all $a_{-i} \in A_{-i}$. Since $v_i^{r_i}$ is strictly increasing, the mapping f is a bijection between $(x_i(a_i, a_{-i}))_{a_{-i} \in A_{-i}} \in \mathbb{R}^{|A_{-i}|}$ and $(y_i(a_{-i}))_{a_{-i} \in A_{-i}} \in Y_i^{|A_{-i}|}$. Also, we have

$$y_i(\pi_i(1)) > y_i(\pi_i(2)) > \cdots > y_i(\pi_i(|A_{-i}|)).$$

Suppose we could show that the set $f(S)$ is a null set with respect to the Lebesgue measure on $Y_i^{|A_{-i}|}$. Since the Lebesgue measure on $Y_i^{|A_{-i}|}$ is the completion of $(\hat{\lambda}_i)^{|A_{-i}|}$, this would imply that there exists a subset S^* such that $f(S) \subset S^* \subset Y_i^{|A_{-i}|}$ and $(\hat{\lambda}_i)^{|A_{-i}|}(S^*) = 0$. Since $\lambda_i^* \ll \hat{\lambda}_i$, we have $(\lambda_i^*)^{|A_{-i}|} \ll (\hat{\lambda}_i)^{|A_{-i}|}$ and hence we would have $(\lambda_i^*)^{|A_{-i}|}(S^*) = 0$. Since λ_i^* is the push forward of the Lebesgue measure λ_i under $v_i^{r_i}$, we would have $(\lambda_i)^{|A_{-i}|}(f^{-i}(S^*)) = 0$, and hence S is a null set with respect to the Lebesgue measure on $\mathbb{R}^{|A_{-i}|}$.

We will now show that the set $f(S)$ is a null set with respect to the Lebesgue measure on $Y_i^{|A_{-i}|}$. The vector $(y_i(a_{-i}))_{a_{-i} \in A_{-i}}$ is completely determined by choosing each of the following:

- (i) a permutation $(\pi_i(1), \pi_i(2), \dots, \pi_i(|A_{-i}|))$ of A_{-i} ,
- (ii) the differences $y_i(\pi_i(t)) - y_i(\pi_i(t+1)) > 0$ for all $1 \leq t < |A_{-i}|$,
- (iii) $y_i(\pi_i(|A_{-i}|)) \in Y_i$ such that

$$y_i(\pi_i(1)) = y_i(\pi_i(|A_{-i}|)) + \sum_{t=1}^{|A_{-i}|-1} y_i(\pi_i(t)) - y_i(\pi_i(t+1)) \in Y_i.$$

Consider the product measure of the following:

- (1) the uniform distribution on the set of permutations of A_{-i} ,

- (2) Lebesgue measure on $y_i(\pi_i(t)) - y_i(\pi_i(t+1)) > 0$ for all $1 \leq t < |A_{-i}|$,
- (3) Lebesgue measure on $y_i(\pi_i(|A_{-i}|)) \in \mathbb{R}$, restricted to $y_i(\pi_i(|A_{-i}|))$ belonging to the interval such that $y_i(\pi_i(|A_{-i}|)) \in Y_i$ and

$$y_i(\pi_i(1)) = y_i(\pi_i(|A_{-i}|)) + \sum_{t=1}^{|A_{-i}|-1} [y_i(\pi_i(t)) - y_i(\pi_i(t+1))] \in Y_i.$$

Now take the push forward of this product measure to the space $Y_i^{|A_{-i}|}$ with respect to the mapping described above. Then, we observe that the completion of this measure is given by the Lebesgue measure on $Y_i^{|A_{-i}|}$.

We will now show that for any fixed permutation $(\pi_i(1), \pi_i(2), \dots, \pi_i(|A_{-i}|))$ and any fixed positive differences $y_i(\pi_i(t)) - y_i(\pi_i(t+1)) > 0$ for all $1 \leq t < |A_{-i}|$, the set of all $y_i(\pi_i(|A_{-i}|))$ such that $(y_i(a_{-i}))_{a_{-i} \in A_{-i}} \in f(S)$ is a null set with respect to the one-dimensional Lebesgue measure.

Let $(\underline{\delta}, \bar{\delta})$ be the largest open interval such that if $y_i(\pi_i(|A_{-i}|)) = \delta$ for any $\delta \in (\underline{\delta}, \bar{\delta})$, then $y_i(\pi_i(|A_{-i}|)), y_i(\pi_i(1)) \in Y_i$. Note that the interval $(\underline{\delta}, \bar{\delta})$ could be empty depending on the fixed positive differences $y_i(\pi_i(t)) - y_i(\pi_i(t+1)) > 0$ for all $1 \leq t < |A_{-i}|$. For $\delta \in (\underline{\delta}, \bar{\delta})$, let Γ^δ denote the game defined by letting $y_i(\pi_i(|A_{-i}|)) := \delta$. In particular, for the game Γ^δ , the payoffs corresponding to player i and action a_i are given by

$$x_i^\delta(a_i, a_{-i}) := (v_i^{r_i})^{-1}(y_i(a_{-i})),$$

for all $a_{-i} \in A_{-i}$, where

$$y_i(a_i) = \delta + \sum_{t=\pi_i^{-1}(a_{-i})}^{|A_{-i}|-1} [y_i(\pi_i(t)) - y_i(\pi_i(t+1))].$$

Consider the function $G_i^{a_i} : \Delta(A_{-i}) \times (\underline{\delta}, \bar{\delta}) \rightarrow \mathbb{R}$, given by

$$G_i^{a_i}(\mu_{-i}, \delta) := \max_{\tilde{a}_i \neq a_i} \mathcal{R}_i[\{(\mu_{-i}(a_{-i}), x_i(\tilde{a}_i, a_{-i}), x_i^\delta(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}}],$$

where the regret function $\mathcal{R}_i[\cdot]$ is as defined in equation (5.5.2). Since the probability weighting functions and the value function for player i are assumed to be continuous, the CPT value function $V_i(L)$ is continuous with respect to the probabilities and the outcomes in the lottery L . Thus, the regret function $\mathcal{R}_i[\cdot]$ is continuous in its arguments, and hence we get that the function $G_i^{a_i}$ is continuous in its arguments.

Now observe that, for any fixed $\delta \in (\underline{\delta}, \bar{\delta})$, the outcomes $(x_i^\delta(a_i, a_{-i}))_{a_{-i} \in A_{-i}}$ are divided into gains and losses depending on the reference point r_i . Hence, for some $0 \leq t_r \leq |A_{-i}|$, we have the outcomes $x_i^\delta(a_i, \pi_i(t)), \forall t \leq t_r$, as gains, and the outcomes $x_i^\delta(a_i, \pi_i(t)), \forall t > t_r$, as losses, where $t_r = 0$ corresponds to the case where all the outcomes $(x_i^\delta(a_i, a_{-i}))_{a_{-i} \in A_{-i}}$

are losses. As a result, the interval $(\underline{\delta}, \bar{\delta})$ can be partitioned into sub-intervals $(\underline{\delta}, \delta_1)$, $[\delta_1, \delta_2), \dots, [\delta_s, \bar{\delta})$, where $\underline{\delta} < \delta_1 < \delta_2 \cdots < \delta_s < \bar{\delta}$, such that over any subinterval I the outcomes are divided into gains and losses at the same point t_r . Here $0 \leq s \leq |A_{-i}|$, with the case $s = 0$ corresponding to the scenario where the division of the outcomes $(x_i^\delta(a_i, a_{-i}))_{a_{-i} \in A_{-i}}$ into gains and losses is the same throughout $(\underline{\delta}, \bar{\delta})$. Note that such an interval I could be open or half-open and half-closed. In the following argument it will not matter whether the subinterval is open or half-open and half closed.

Let us now consider the function $G_i^{a_i}$ restricted to $\Delta(A_{-i}) \times I$ for a fixed subinterval I . Let $0 \leq t_r \leq |A_{-i}|$ be the point that divides the outcomes $(x_i^\delta(a_i, a_{-i}))_{a_{-i} \in A_{-i}}$ into gains and losses. Suppose we could show that the set of $\delta \in I$ such that $(y_i(a_{-i}))_{a_{-i} \in A_{-i}} \in f(S)$ is a null set with respect to the one-dimensional Lebesgue measure. Since this would be true for each of the subintervals I , and there are only finitely many such subintervals in the partitioning of $(\underline{\delta}, \bar{\delta})$ above, we would get the desired result.

We first prove the following useful property: For any $\delta, \tilde{\delta} \in I$, and $\mu_{-i} \in \Delta(A_{-i})$, we have

$$G_i^{a_i}(\mu_{-i}, \delta) - G_i^{a_i}(\mu_{-i}, \tilde{\delta}) = W_i(\mu_{-i})(\tilde{\delta} - \delta), \quad (5.H.1)$$

where

$$W_i(\mu_{-i}) := w_i^+ \left(\sum_{t=1}^{t_r} \mu_{-i}(\pi_i(t)) \right) + w_i^- \left(\sum_{t=t_r+1}^{|A_{-i}|} \mu_{-i}(\pi_i(t)) \right).$$

To see this, write

$$\begin{aligned} G_i^{a_i}(\mu_{-i}, \delta) &= \left(\max_{\tilde{a}_i \neq a_i} V_i(\{(\mu_{-i}(a_{-i}), x_i(\tilde{a}_i, a_{-i}))\}_{a_{-i} \in A_{-i}}) \right) \\ &\quad - V_i(\{(\mu_{-i}(a_{-i}), x_i^\delta(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}}), \end{aligned}$$

which gives

$$\begin{aligned} G_i^{a_i}(\mu_{-i}, \delta) - G_i^{a_i}(\mu_{-i}, \tilde{\delta}) &= V_i(\{(\mu_{-i}(a_{-i}), x_i^{\tilde{\delta}}(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}}) \\ &\quad - V_i(\{(\mu_{-i}(a_{-i}), x_i^\delta(a_i, a_{-i}))\}_{a_{-i} \in A_{-i}}). \end{aligned}$$

Equation (5.H.1) then follows from equation (1.3.9).

Note that $W_i(\mu_{-i}) > 0$ always. Indeed, since

$$\sum_{t=1}^{t_r} \mu_{-i}(\pi_i(t)) + \sum_{t=t_r+1}^{|A_{-i}|} \mu_{-i}(\pi_i(t)) = 1,$$

at least one of these two summations is positive, and $w_i^\pm(p) > 0$ for $p > 0$ from the assumptions on the probability weighting functions.

For any $\delta \in I$, we have $\mu_{-i} \in C(\Gamma^\delta, i, a_i)$ if and only if $G_i^{a_i}(\mu_{-i}, \delta) \leq 0$. If $G_i^{a_i}(\mu_{-i}, \delta) < 0$ then, by the continuity of the function $G_i^{a_i}$, we will have a neighborhood around the point

μ_{-i} that belongs to $C(\Gamma^\delta, i, a_i)$. Since the domain $\Delta(A_{-i})$ itself does not have any isolated points, it prevents μ_{-i} from being an isolated point of $C(\Gamma^\delta, i, a_i)$. Thus, the fact that μ_{-i} is an isolated point of $C(\Gamma^\delta, i, a_i)$ implies that $G_i^{a_i}(\mu_{-i}, \delta) = 0$. If μ_{-i} is not a strict local minimum of $G_i^{a_i}(\cdot, \delta)$, then there exists a sequence of points $(\mu_{-i}^t)_{t \geq 1}$ converging to μ_{-i} such that $G_i^{a_i}(\mu_{-i}^t, \delta) \leq 0$, for all $t \geq 1$. Then the sequence $(\mu_{-i}^t)_{t \geq 1}$ belongs to the set $C(\Gamma^\delta, i, a_i)$, contradicting the fact that μ_{-i} is an isolated point in the set $C(\Gamma^\delta, i, a_i)$. We have shown that if μ_{-i} is an isolated point in the set $C(\Gamma^\delta, i, a_i)$, this implies that $G_i^{a_i}(\mu_{-i}, \delta) = 0$ and that μ_{-i} is a strict local minimum of $G_i^{a_i}(\tilde{\mu}_{-i}, \delta)$ as a function of $\tilde{\mu}_{-i} \in \Delta(A_{-i})$.

To complete the proof of the proposition, it is enough to show that the set of all $\delta \in I$ for which there exists $\mu_{-i} \in \Delta(A_{-i})$ such that $G_i^{a_i}(\mu_{-i}, \delta) = 0$ and μ_{-i} is a strict local minimum of $G_i^{a_i}(\cdot, \delta)$ is a null set with respect to one dimensional Lebesgue measure. Let $T \subset \Delta(A_{-i}) \times I$ be the set of all pairs (μ_{-i}, δ) such that $G_i^{a_i}(\mu_{-i}, \delta) = 0$ and μ_{-i} is a strict local minimum of $G_i^{a_i}(\cdot, \delta)$. We will prove that the set T is countable. To see this, for each pair $(\mu_{-i}, \delta) \in T$, there exists a pair of vectors with rational elements, $(p_{\mu_{-i}, \delta}(a_{-i}))_{a_{-i} \in A_{-i}}$ and $(q_{\mu_{-i}, \delta}(a_{-i}))_{a_{-i} \in A_{-i}}$, such that

$$p_{\mu_{-i}, \delta}(a_{-i}) < \mu_{-i}(a_{-i}) < q_{\mu_{-i}, \delta}(a_{-i}), \text{ for all } a_{-i} \in A_{-i},$$

and for any $\tilde{\mu}_{-i} \in \Delta(A_{-i})$ such that

$$p_{\mu_{-i}, \delta}(a_{-i}) < \tilde{\mu}_{-i}(a_{-i}) < q_{\mu_{-i}, \delta}(a_{-i}), \text{ for all } a_{-i} \in A_{-i},$$

we have $G_i^{a_i}(\tilde{\mu}_{-i}, \delta) > G_i^{a_i}(\mu_{-i}, \delta)$. Suppose there are two distinct pairs $(\mu'_{-i}, \delta'), (\mu''_{-i}, \delta'') \in T$ such that $p_{\mu'_{-i}, \delta'}(a_{-i}) = p_{\mu''_{-i}, \delta''}(a_{-i}) =: p(a_{-i})$ and $q_{\mu'_{-i}, \delta'}(a_{-i}) = q_{\mu''_{-i}, \delta''}(a_{-i}) =: q(a_{-i})$ for all $a_{-i} \in A_{-i}$. We note that in this case we must have $\delta' \neq \delta''$. Let $\delta' < \delta''$ without loss of generality. We have $G_i^{a_i}(\mu'_{-i}, \delta'') \geq 0$ because

$$p(a_{-i}) < \mu'_{-i}(a_{-i}) < q(a_{-i}), \text{ for all } a_{-i} \in A_{-i}.$$

From equation (5.H.1), we have

$$G_i^{a_i}(\mu'_{-i}, \delta') - G_i^{a_i}(\mu'_{-i}, \delta'') = W_i(\mu'_{-i})(\delta'' - \delta') > 0.$$

This implies $G_i^{a_i}(\mu'_{-i}, \delta') > 0$ contradicting $(\mu'_{-i}, \delta') \in T$. Thus we have an injective map from the set T to the set $\mathbb{Q}^{2|A_{-i}|}$. Hence the set T is countable. Thus the set of all $\delta \in I$, for which there exists a μ_{-i} such that $(\mu_{-i}, \delta) \in T$ is also countable and hence a null set. This completes the proof. \square

5.I Proof of proposition 5.5.1

Since $\Delta(A)$ is a compact set, ξ^t converges to the set $C(\Gamma)$ iff for every convergent subsequence ξ^{t_k} , say, converging to $\hat{\xi}$, we have $\hat{\xi} \in C(\Gamma)$. Let $\xi^{t_k} \rightarrow \hat{\xi}$ be a convergent subsequence. For each player i , and for every $a_i, \tilde{a}_i \in A_i, a_i \neq \tilde{a}_i$ such that $\hat{\xi}_i(a_i) > 0$, we have

$$K_i^{t_k}(a_i, \tilde{a}_i) \rightarrow \hat{\xi}_i(a_i) \mathcal{R}_i \left[\left\{ \left(\hat{\xi}_{-i}(a_{-i} | a_i), x_i(\tilde{a}_i, a_{-i}), x_i(a_i, a_{-i}) \right) \right\}_{a_{-i} \in A_{-i}} \right], \quad (5.I.1)$$

by continuity of $V_i(p, x)$ as a function of the probability vector p for a fixed outcome profile x . The result is immediate from the definition of CPT correlated equilibrium. \square

5.J Proof of Lemma 5.5.4

We will first use the fact that player 2 is randomizing over her actions I and III, independently at all the steps $(t_{odd}^l)_{l \geq 1}$, and show that for sufficiently large l , $v_{odd}^l(0, I)$ and $v_{odd}^l(0, III)$ are almost equal with high probability. To see this, observe that the sequence $(M_l, l \geq 1)$ is a martingale, where

$$M_l := l \times (\nu_{odd}^l(0, I) - \nu_{odd}^l(0, III)).$$

Indeed, letting $M_1^l := (M_1, \dots, M_l)$, we have

$$\begin{aligned} \mathbb{E}[M_{l+1} - M_l | M_1^l] &= \mathbb{E}[M_{l+1} - M_l | M_1^l, a_1^{t_{odd}^{l+1}} = 0] P(a_1^{t_{odd}^{l+1}} = 0 | M_1^l) \\ &\quad + \mathbb{E}[M_{l+1} - M_l | M_1^l, a_1^{t_{odd}^{l+1}} = 1] P(a_1^{t_{odd}^{l+1}} = 1 | M_1^l) \\ &= \mathbb{E}[\mathbf{1}\{a_1^{t_{odd}^{l+1}} = (0, I)\} - \mathbf{1}\{a_1^{t_{odd}^{l+1}} = (0, III)\} | M_1^l, a_1^{t_{odd}^{l+1}} = 0] P(a_1^{t_{odd}^{l+1}} = 0 | M_1^l) + 0 \\ &= \frac{1}{2} - \frac{1}{2} = 0, \end{aligned}$$

where the last line follows from the fact that player 2 plays σ_{odd} at each of the steps t_{odd}^l independently. Thus, for example by the Azuma-Hoeffding inequality, for any $\delta > 0$, there exists an integer $l_\delta^{(1)} > 1$, such that for all $l \geq l_\delta^{(1)}$, equation (5.5.7) holds. Similarly, there exist integers $l_\delta^{(2)}, l_\delta^{(3)}, l_\delta^{(4)} > 1$, such that for all $l \geq l_\delta^{(2)}$, equation (5.5.8) holds, for all $l \geq l_\delta^{(3)}$, equation (5.5.9) holds, and for all $l \geq l_\delta^{(4)}$, equation (5.5.10) holds. This taking

$$l_\delta := \max\{l_\delta^{(1)}, l_\delta^{(2)}, l_\delta^{(3)}, l_\delta^{(4)}\},$$

we get the required result. \square

5.K Proof of Proposition 5.5.3

Here are two simple technical lemmas that we will use repeatedly in the rest of the discussion in this section. The proof of each of these lemmas is elementary, and is therefore omitted.

Lemma 5.K.1. *If $P(F_1) > \alpha$ and $P(F_2) \geq \beta$ such that $\alpha + \beta > 1$, then*

$$P(F_1 | F_2) \geq P(F_1 \cap F_2) > \alpha - (1 - \beta).$$

\square

Lemma 5.K.2. *If $\delta > 0$, and x, y, a, b are real numbers such that $x+y \in [a, b]$ and $|x-y| < \delta$, then*

$$x, y \in ((a - \delta)/2, (b + \delta)/2).$$

□

Let E_1^k denote the event that the following inclusion holds:

$$\xi^{T^{k+1}}(1, \cdot) \in \left[\left(\frac{1 - f_3^{k+1}}{2}, 0, \frac{1 - f_3^{k+1}}{2}, 0 \right) \right]_{\delta_1}. \quad (5.K.1)$$

Lemma 5.K.3. *Recall that k_0 is defined in equation (5.5.17). For any $k \geq k_0$, if $P(f_3^{k+1} < 1 - \delta_2) > 1/4$, then*

$$P(E_1^k | f_3^{k+1} < 1 - \delta_2) > 1/4 - \delta_1. \quad (5.K.2)$$

Proof. Fix $k \geq k_0$. From inequality (5.5.18) and the assumption $T > 2/\delta_1$, we have

$$\xi^{T^{k+1}}(1, \text{II}) + \xi^{T^{k+1}}(1, \text{IV}) < \frac{2}{T} < \delta_1, \quad (5.K.3)$$

and hence each term is strictly less than δ_1 , i.e.

$$\xi^{T^{k+1}}(1, \text{II}), \xi^{T^{k+1}}(1, \text{IV}) \in [0, \delta_1). \quad (5.K.4)$$

Since $k \geq k_0$, from (5.5.17) and (5.5.8), for $l := \max\{l : t_{\text{odd}}^l \leq T^{k+1}\}$, we have

$$\begin{aligned} & P\left(|\xi^{T^{k+1}}(1, \text{I}) - \xi^{T^{k+1}}(1, \text{III})| < \delta_1\right) \\ &= P\left(|\nu_{\text{odd}}^l(1, \text{I}) - \nu_{\text{odd}}^l(1, \text{III})|(1 - f_1^{k+1}) < \delta_1\right) \\ &\geq P\left(|\nu_{\text{odd}}^l(1, \text{I}) - \nu_{\text{odd}}^l(1, \text{III})| < \delta_1\right) > 1 - \delta_1, \end{aligned}$$

In Lemma 5.K.1, taking F_1 to be the event

$$\left\{ |\xi^{T^{k+1}}(1, \text{I}) - \xi^{T^{k+1}}(1, \text{III})| < \delta_1 \right\},$$

and F_2 to be the event $\{f_3^{k+1} < 1 - \delta_2\}$, we have $P(F_1) > 1 - \delta_1$, $P(F_2) > 1/4$. Since $\delta_1 < 1/4$, we have

$$P\left(|\xi^{T^{k+1}}(1, \text{I}) - \xi^{T^{k+1}}(1, \text{III})| < \delta_1 \mid f_3^{k+1} < 1 - \delta_2\right) > 1/4 - \delta_1. \quad (5.K.5)$$

Since

$$\xi^{T^{k+1}}(1, \text{I}) + \xi^{T^{k+1}}(1, \text{II}) + \xi^{T^{k+1}}(1, \text{III}) + \xi^{T^{k+1}}(1, \text{IV}) = 1 - f_3^{k+1},$$

combined with (5.K.3), we have

$$\xi^{T^{k+1}}(1, \text{I}) + \xi^{T^{k+1}}(1, \text{III}) \in [1 - f_3^{k+1} - \delta_1, 1 - f_3^{k+1}]. \quad (5.K.6)$$

From (5.K.5), (5.K.6) and Lemma 5.K.2, we have

$$P\left(\xi^{T^{k+1}}(1, I), \xi^{T^{k+1}}(1, III) \in \left(\frac{1-f_3^{k+1}-\delta_1}{2} - \frac{\delta_1}{2}, \frac{1-f_3^{k+1}}{2} + \frac{\delta_1}{2}\right) \middle| f_3^{k+1} < 1 - \delta_2\right) > 1/4 - \delta_1.$$

Combined with (5.K.4), we get (5.K.2) and this completes the proof of the lemma. \square

Lemma 5.K.4. *For any $k \geq k_0$, if $P(f_3^{k+1} < 1 - \delta_2) > 1/4$, then*

$$P\left([K_1^{T^{k+1}}(1, 0)]^+ > \delta_2 c_1\right) > \frac{1}{4} \left(\frac{1}{4} - \delta_1\right). \quad (5.K.7)$$

Proof. From (5.5.18), we know that player 2 plays σ_{odd} for at least a fraction $1 - \frac{2}{T}$ of the steps up to step $t = T^{k+1}$. Since action 1 is not a best response of player 1 for σ_{odd} , we will now show that, if player 1 does not play action 0 for a sufficiently high fraction of steps up to step $t = T^{k+1}$, then she will have a significant regret $K_1^{T^{k+1}}(1, 0)$. More precisely, for any $k \geq k_0$, if $f_3^{k+1} < 1 - \delta_2$ and the inclusion (5.K.1) holds, then we can write

$$\begin{aligned} \xi_{-1}^{T^{k+1}}(\cdot|1)\xi_1^{T^{k+1}}(1) &\in \left[\left(\frac{1-f_3^{k+1}}{2}, 0, \frac{1-f_3^{k+1}}{2}, 0\right)\right]_{\delta_1} \\ \iff \xi_{-1}^{T^{k+1}}(\cdot|1)(1-f_3^{k+1}) &\in \left[\left(\frac{1-f_3^{k+1}}{2}, 0, \frac{1-f_3^{k+1}}{2}, 0\right)\right]_{\delta_1} \\ \iff \xi_{-1}^{T^{k+1}}(\cdot|1) &\in \left[\left(\frac{1}{2}, 0, \frac{1}{2}, 0\right)\right]_{\delta_1/(1-f_3^{k+1})} \\ \implies \xi_{-1}^{T^{k+1}}(\cdot|1) &\in [\sigma_{odd}]_{\frac{\delta_1}{\delta_2}}. \end{aligned}$$

Hence, from (5.5.16) and (5.5.15), we have

$$K_1^{T^{k+1}}(1, 0) = \xi_1^{T^{k+1}}(1)\mathcal{R}_1 \left[\left\{ \left(\xi_{-1}^{T^{k+1}}(\cdot|1), x_1(0, \cdot), x_1(1, \cdot) \right) \right\} \right] > \delta_2 c_1,$$

on the event where $f_3^{k+1} < 1 - \delta_2$ and the inclusion (5.K.1) holds. Thus for any $k \geq k_0$, if $P(f_3^{k+1} < 1 - \delta_2) > 1/4$, then we have

$$\begin{aligned} &P\left([K_1^{T^{k+1}}(1, 0)]^+ > \delta_2 c_1\right) \\ &= P\left([K_1^{T^{k+1}}(1, 0)]^+ > \delta_2 c_1 \middle| f_3^{k+1} < 1 - \delta_2\right) P(f_3^{k+1} < 1 - \delta_2) \\ &\geq P(E_1^k | f_3^{k+1} < 1 - \delta_2) P(f_3^{k+1} < 1 - \delta_2) \\ &> \frac{1}{4} \left(\frac{1}{4} - \delta_1\right), \end{aligned}$$

where the last but one inequality follows from the fact that E_1^k and $\{f_3^{k+1} < 1 - \delta_2\}$ imply $[K_1^{T^{k+1}}(1, 0)]^+ > \delta_2 c_1$, and the last inequality follows from the condition $P(f_3^{k+1} < 1 - \delta_2) > 1/4$ and Lemma 5.K.3. \square

	I	II	III	IV
0	0.25	$0.25f_4^{k+1}$	0.25	$0.25f_4^{k+1}$
1	0	$0.25(1 - f_4^{k+1})$	0	$0.25(1 - f_4^{k+1})$

Table 5.K.1: Empirical distribution $\hat{\mu}$ in Example 5.5.2.

Consider now the probability distribution $\hat{\mu}$ shown in Table 5.K.1. Recall that f_4^{k+1} is the fraction of times player 1 plays action 0 among the steps from step $T^{k+1} + 1$ to step $2T^{k+1}$. Note that, since f_4^{k+1} is a random variable, so is $\hat{\mu}$.

Lemma 5.K.5. *For all $k \geq k_0$, if $P(f_3^{k+1} < 1 - \delta_2) \leq 1/4$, then*

$$P(\xi^{2T^{k+1}} \in [\hat{\mu}]_{\delta_2} | f_3^{k+1} \geq 1 - \delta_2) > 1/4 - 3\delta_1. \quad (5.K.8)$$

We also recall that $\delta_1 < 1/16$, so the lower bound in (5.K.8) is strictly positive.

Proof. Since player 2 plays σ_{even} from step $T^{k+1} + 1$ to step $2T^{k+1}$, if $f_3^{k+1} \geq 1 - \delta_2$, then

$$\xi^{2T^{k+1}}(1, \text{I}) + \xi^{2T^{k+1}}(1, \text{III}) \leq \xi_1^{T^{k+1}}(1)/2 = (1 - f_3^{k+1})/2 \leq \delta_2/2. \quad (5.K.9)$$

This means that each term is strictly less than δ_2 , so we have

$$\xi^{2T^{k+1}}(1, \text{I}), \xi^{2T^{k+1}}(1, \text{III}) \in [0, \delta_2). \quad (5.K.10)$$

Further, from equation (5.5.19) and the assumption $T > 2/\delta_1$, we have

$$\begin{aligned} \xi^{2T^{k+1}}(0, \text{I}) + \xi^{2T^{k+1}}(0, \text{III}) + \xi^{2T^{k+1}}(1, \text{I}) + \xi^{2T^{k+1}}(1, \text{III}) \\ = 1 - f_2^{k+1} \in [0.5 - 1/T, 0.5] \subset [0.5 - \delta_1, 0.5]. \end{aligned}$$

Combining this with (5.K.9), we have

$$\xi^{2T^{k+1}}(0, \text{I}) + \xi^{2T^{k+1}}(0, \text{III}) \in [0.5 - \delta_1 - \delta_2/2, 0.5], \quad (5.K.11)$$

on the event where $f_3^{k+1} \geq 1 - \delta_2$. Since $k \geq k_0$, from (5.5.17) and (5.5.7), for $l := \max\{l : t_{\text{odd}}^l \leq 2T^{k+1}\}$, we have

$$\begin{aligned} P\left(|\xi^{2T^{k+1}}(0, \text{I}) - \xi^{2T^{k+1}}(0, \text{III})| < \delta_1\right) \\ = P\left(|\nu_{\text{odd}}^l(0, \text{I}) - \nu_{\text{odd}}^l(0, \text{III})|(1 - f_2^{k+1}) < \delta_1\right) \\ \geq P\left(|\nu_{\text{odd}}^l(0, \text{I}) - \nu_{\text{odd}}^l(0, \text{III})| < \delta_1\right) > 1 - \delta_1. \end{aligned}$$

In Lemma 5.K.1, taking F_1 to be the event

$$\left\{|\xi^{2T^{k+1}}(0, \text{I}) - \xi^{2T^{k+1}}(0, \text{III})| < \delta_1\right\}$$

and F_2 to be the event $\{f_3^{k+1} \geq 1 - \delta_2\}$, we have $P(F_1) > 1 - \delta_1$, $P(F_2) \geq 3/4$. Since $\delta_1 < 1/4$, we have

$$P\left(|\xi^{2T^{k+1}}(0,\text{I}) - \xi^{2T^{k+1}}(0,\text{III})| < \delta_1 \mid f_3^{k+1} \geq 1 - \delta_2\right) > 3/4 - \delta_1. \quad (5.K.12)$$

Form (5.K.11), (5.K.12) and Lemma 5.K.2, we have

$$\begin{aligned} P\left(\xi^{2T^{k+1}}(0,\text{I}), \xi^{2T^{k+1}}(0,\text{III}) \in (0.25 - \delta_1 - \delta_2/4, 0.25 + \delta_1/2) \mid f_3^{k+1} \geq 1 - \delta_2\right) \\ > 3/4 - \delta_1. \end{aligned}$$

Here we note that $0.25 - \delta_1 - \delta_2/4 > 0$. Since $\epsilon_1 < 0.5$ and $\delta_1 = \epsilon_1\delta_2$, we have

$$P\left(\xi^{2T^{k+1}}(0,\text{I}), \xi^{2T^{k+1}}(0,\text{III}) \in (0.25 - \delta_2, 0.25 + \delta_2) \mid f_3^{k+1} \geq 1 - \delta_2\right) > 3/4 - \delta_1. \quad (5.K.13)$$

From (5.5.18) and the assumption $T > 2/\delta_1$, we have

$$\begin{aligned} \xi^{2T^{k+1}}(0,\text{II}) + \xi^{2T^{k+1}}(0,\text{IV}) &\in [0.5f_4^{k+1}, 0.5f_4^{k+1} + 0.5f_1^{k+1}] \\ &\in [0.5f_4^{k+1}, 0.5f_4^{k+1} + \delta_1]. \end{aligned} \quad (5.K.14)$$

Since $k \geq k_0$, from (5.5.17) and (5.5.9), for $l := \max\{l : t_{even}^l \leq 2T^{k+1}\}$, we have

$$\begin{aligned} P\left(|\xi^{2T^{k+1}}(0,\text{II}) - \xi^{2T^{k+1}}(0,\text{IV})| < \delta_1\right) \\ = P\left(|\nu_{even}^l(0,\text{II}) - \nu_{even}^l(0,\text{IV})|(f_2^{k+1}) < \delta_1\right) \\ \geq P\left(|\nu_{even}^l(0,\text{II}) - \nu_{even}^l(0,\text{IV})| < \delta_1\right) > 1 - \delta_1. \end{aligned}$$

In Lemma 5.K.1, taking F_1 to be the event

$$\left\{|\xi^{2T^{k+1}}(0,\text{II}) - \xi^{2T^{k+1}}(0,\text{IV})| < \delta_1\right\}$$

and F_2 to be the event $\{f_3^{k+1} \geq 1 - \delta_2\}$, we have $P(F_1) > 1 - \delta_1$, $P(F_2) \geq 3/4$. Since $\delta_1 < 1/4$, we have

$$P\left(|\xi^{2T^{k+1}}(0,\text{II}) - \xi^{2T^{k+1}}(0,\text{IV})| < \delta_1 \mid f_3^{k+1} \geq 1 - \delta_2\right) > 3/4 - \delta_1. \quad (5.K.15)$$

From (5.K.14), (5.K.15) and Lemma 5.K.2, we have

$$\begin{aligned} P\left(\xi^{2T^{k+1}}(0,\text{II}), \xi^{2T^{k+1}}(0,\text{IV}) \in (0.25f_4^{k+1} - \delta_1, 0.25f_4^{k+1} + \delta_1) \mid f_3^{k+1} \geq 1 - \delta_2\right) \\ > 3/4 - \delta_1, \end{aligned} \quad (5.K.16)$$

Note that here $0.25f_4^{k+1} - \delta_1$ could be negative. From (5.5.18) and the assumption $T > 2/\delta_1$, we have

$$\begin{aligned} \xi^{2T^{k+1}}(1, \text{II}) + \xi^{2T^{k+1}}(1, \text{IV}) &\in [0.5(1 - f_4^{k+1}), 0.5(1 - f_4^{k+1}) + 0.5f_1^{k+1}] \\ &\in [0.5(1 - f_4^{k+1}), 0.5(1 - f_4^{k+1}) + \delta_1]. \end{aligned} \quad (5.K.17)$$

Since $k \geq k_0$, from (5.5.17) and (5.5.10), for $l := \max\{l : t_{\text{even}}^l \leq 2T^{k+1}\}$, we have

$$\begin{aligned} &P\left(|\xi^{2T^{k+1}}(1, \text{II}) - \xi^{2T^{k+1}}(1, \text{IV})| < \delta_1\right) \\ &= P\left(|\nu_{\text{even}}^l(1, \text{II}) - \nu_{\text{even}}^l(1, \text{IV})|(f_2^{k+1}) < \delta_1\right) \\ &\geq P\left(|\nu_{\text{even}}^l(1, \text{II}) - \nu_{\text{even}}^l(1, \text{IV})| < \delta_1\right) > 1 - \delta_1. \end{aligned}$$

In Lemma 5.K.1, taking F_1 to be the event

$$\left\{|\xi^{2T^{k+1}}(1, \text{II}) - \xi^{2T^{k+1}}(1, \text{IV})| < \delta_1\right\}$$

and F_2 to be the event $\{f_3^{k+1} \geq 1 - \delta_2\}$, we have $P(F_1) > 1 - \delta_1$, $P(F_2) \geq 3/4$. Since $\delta_1 < 1/4$, we have

$$P\left(|\xi^{2T^{k+1}}(1, \text{II}) - \xi^{2T^{k+1}}(1, \text{IV})| < \delta_1 \mid f_3^{k+1} \geq 1 - \delta_2\right) > 3/4 - \delta_1. \quad (5.K.18)$$

Form (5.K.17), (5.K.18) and Lemma 5.K.2, we have

$$\begin{aligned} &P\left(\xi^{2T^{k+1}}(1, \text{II}), \xi^{2T^{k+1}}(1, \text{IV}) \in (0.25(1 - f_4^{k+1}) - \delta_1, 0.25(1 - f_4^{k+1}) + \delta_1) \mid f_3^{k+1} \geq 1 - \delta_2\right) \\ &> 3/4 - \delta_1. \end{aligned} \quad (5.K.19)$$

Note that $0.25(1 - f_4^{k+1}) - \delta_1$ could be negative. From (5.K.13), (5.K.16), (5.K.19) and (5.K.10) we get (5.K.8), and this completes the proof. \square

We now consider two scenarios based on whether $f_4^{k+1} < 1 - \delta_3$ or $f_4^{k+1} \geq 1 - \delta_3$.

Lemma 5.K.6. *For any $k \geq k_0$, if $f_4^{k+1} < 1 - \delta_3$ and $\xi^{2T^{k+1}} \in [\hat{\mu}]_{\delta_2}$, then $K_1^{2T^{k+1}}(1, 0) > 0.5\delta_3c_2$.*

Proof. If $f_4^{k+1} < 1 - \delta_3$, then $\xi^{2T^{k+1}} \in [\hat{\mu}]_{\delta_2}$ implies that $\xi_{-1}^{2T^{k+1}}(\cdot|1) \in [\sigma_{\text{even}}]_{(2\delta_2)/(0.5\delta_3)}$. Indeed, since $\xi_1^{2T^{k+1}}(1) \geq (1 - f_4^{k+1})/2 > 0.5\delta_3$, normalizing $\xi^{2T^{k+1}}(1, \cdot)$ by $\xi_1^{2T^{k+1}}(1)$, we get

$$\xi_{-1}^{2T^{k+1}}(\text{I}|1), \xi_{-1}^{2T^{k+1}}(\text{III}|1) \in [0, 2\delta_2/\delta_3], \quad (5.K.20)$$

and

$$|\xi_{-1}^{2T^{k+1}}(\text{II}|1) - \xi_{-1}^{2T^{k+1}}(\text{IV}|1)| < \frac{4\delta_2}{\delta_3}. \quad (5.K.21)$$

Since

$$\xi_{-1}^{2T^{k+1}}(\text{I}|1) + \xi_{-1}^{2T^{k+1}}(\text{II}|1) + \xi_{-1}^{2T^{k+1}}(\text{III}|1) + \xi_{-1}^{2T^{k+1}}(\text{IV}|1) = 1,$$

we have,

$$\xi_{-1}^{2T^{k+1}}(\text{II}|1) + \xi_{-1}^{2T^{k+1}}(\text{IV}|1) \in \left[1 - \frac{4\delta_2}{\delta_3}, 1\right]. \quad (5.K.22)$$

From (5.K.21), (5.K.22) and Lemma 5.K.2, we have

$$\xi_{-1}^{2T^{k+1}}(\text{II}|1), \xi_{-1}^{2T^{k+1}}(\text{IV}|1) \in \left(\frac{1}{2} - \frac{4\delta_2}{\delta_3}, \frac{1}{2} + \frac{2\delta_2}{\delta_3}\right), \quad (5.K.23)$$

and hence $\xi_{-1}^{2T^{k+1}}(\cdot|1) \in [\sigma_{\text{even}}]_{(4\delta_2)/\delta_3}$. Then, from the assumption (5.5.14) we have $\xi_{-1}^{2T^{k+1}}(\cdot|1) \in [\sigma_{\text{even}}]_{\epsilon_2}$, and hence from (5.5.13) we have

$$K_1^{2T^{k+1}}(1,0) = \xi_1^{2T^{k+1}}(1)\mathcal{R}_1 \left[\left\{ \left(\xi_{-1}^{2T^{k+1}}(\cdot|1), x_1(0, \cdot), x_1(1, \cdot) \right) \right\} \right] > 0.5\delta_3 c_2. \quad (5.K.24)$$

□

Lemma 5.K.7. *For any $k \geq k_0$, if $f_4^{k+1} \geq 1 - \delta_3$, $f_3^{k+1} \geq 1 - \delta_2$ and $\xi^{2T^{k+1}} \in [\hat{\mu}]_{\delta_2}$, then $K_1^{2T^{k+1}}(0,1) > (1 - \delta_3)c_3$.*

Proof. If $f_4^{k+1} \geq 1 - \delta_3$ and $f_3^{k+1} \geq 1 - \delta_2$, then $\xi^{2T^{k+1}} \in [\hat{\mu}]_{\delta_2}$ implies that

$$\xi_{-1}^{2T^{k+1}}(\cdot|0) \in [\sigma_{\text{unif}}]_{\frac{\delta_3/4 + \delta_2 + \delta_3/8 + \delta_2/8}{1 - \delta_3/2 - \delta_2/2}}. \quad (5.K.25)$$

To see this, note that $f_4^{k+1} \geq 1 - \delta_3$ and $\xi^{2T^{k+1}} \in [\hat{\mu}]_{\delta_2}$ imply that $\xi^{2T^{k+1}}(0, \cdot) \in [\sigma_{\text{unif}}]_{\delta_3/4 + \delta_2}$. We have $\xi_1^{2T^{k+1}}(0) = f_3^{k+1}/2 + f_4^{k+1}/2 \geq 1 - \delta_3/2 - \delta_2/2$. Let $\kappa := (1 - \xi_1^{2T^{k+1}}(0))/4$. Thus $0 \leq \kappa \leq \delta_3/8 + \delta_2/8$. Let $\sigma_{\text{unif}} - \kappa := (0.25 - \kappa, 0.25 - \kappa, 0.25 - \kappa, 0.25 - \kappa)$. Then we have

$$\xi^{2T^{k+1}}(0, \cdot) \in [0.25 - \kappa, 0.25 - \kappa, 0.25 - \kappa, 0.25 - \kappa]_{\delta_3/4 + \delta_2 + \kappa}.$$

Normalizing $\sigma_{\text{unif}} - \kappa$ with $1 - 4\kappa = \xi_1^{2T^{k+1}}(0)$ gives us σ_{unif} . As a result, normalizing $\xi^{2T^{k+1}}(0, \cdot)$ with $\xi_1^{2T^{k+1}}(0)$ gives (5.K.25). Then, from the assumptions $\epsilon_3 < 1$, $\epsilon_2 < 1$, $\delta_2 = \epsilon_2\delta_3/4$ and $\delta_3 = \epsilon_3/2$, we have

$$\frac{\delta_3/4 + \delta_2 + \delta_3/8 + \delta_2/8}{1 - \delta_3/2 - \delta_2/2} \leq \frac{\delta_3}{1 - \delta_3} \leq \epsilon_3.$$

Thus, $\xi_{-1}^{2T^{k+1}}(\cdot|0) \in [\sigma_{\text{unif}}]_{\epsilon_3}$, and hence from (5.5.11) we have

$$K_1^{2T^{k+1}}(0,1) = \xi_1^{2T^{k+1}}(0)\mathcal{R}_1 \left[\left\{ \left(\xi_{-1}^{2T^{k+1}}(\cdot|0), x_1(1, \cdot), x_1(0, \cdot) \right) \right\} \right] \quad (5.K.26)$$

$$> (1 - \delta_3/2 - \delta_2/2)c_3 > (1 - \delta_3)c_3. \quad (5.K.27)$$

□

Lemma 5.K.8. For any $k \geq k_0$, if $P(f_3^{k+1} < 1 - \delta_2) \leq 1/4$, then

$$P(\bar{K}^k > \min\{0.5\delta_3c_2, (1 - \delta_3)c_3\}) > \frac{3}{4} \left(\frac{1}{4} - 3\delta_1 \right), \quad (5.K.28)$$

where \bar{K}^k is defined in equation (5.5.4).

Proof. From Lemma 5.K.6 and Lemma 5.K.7 we obtain the following: if $f_3^{k+1} \geq 1 - \delta_2$ and $\xi^{2T^{k+1}} \in [\hat{\mu}]_{\delta_2}$, then $\bar{K}^k > \min\{0.5\delta_3c_2, (1 - \delta_3)c_3\}$. As a result, from Lemma 5.K.5, if $P(f_3^{k+1} < 1 - \delta_2) \leq 1/4$, then

$$\begin{aligned} & P(\bar{K}^k > \min\{0.5\delta_3c_2, (1 - \delta_3)c_3\}) \\ & \geq P(\bar{K}^k > \min\{0.5\delta_3c_2, (1 - \delta_3)c_3\} | f_3^{k+1} \geq 1 - \delta_2) P(f_3^{k+1} \geq 1 - \delta_2) \\ & \geq P(\xi^{2T^{k+1}} \in [\hat{\mu}]_{\delta_2} | f_3^{k+1} \geq 1 - \delta_2) P(f_3^{k+1} \geq 1 - \delta_2) \\ & > \frac{3}{4} \left(\frac{1}{4} - 3\delta_1 \right). \end{aligned}$$

□

Proof of Proposition 5.5.3. Take

$$\tilde{\epsilon} = \min\{\delta_2c_1, 0.5\delta_3c_2, (1 - \delta_3)c_3\}$$

and

$$\tilde{\delta} = \min\left\{ \frac{1}{4} \left(\frac{1}{4} - \delta_1 \right), \frac{3}{4} \left(\frac{1}{4} - 3\delta_1 \right) \right\}.$$

From Lemma 5.K.4 and Lemma 5.K.8 it follows that for all $k \geq k_0$,

$$P(\bar{K}^k > \tilde{\epsilon}) > \tilde{\delta},$$

and this concludes the proof. □

Notes

¹⁵Also known as the internal regret or the conditional regret.

¹⁶Foster and Vohra [50] refer to it as the best response. In order to avoid confusion with the best response set defined in Chapter 3, we prefer to use the term best reaction.

¹⁷Foster and Vohra [49] prove the existence of a randomized forecasting scheme that makes the forecaster's calibration score, i.e. the expression in equation (5.4.1), tend to zero in probability. However, as noted in [27], the same argument proves that the convergence of the calibration score holds, in fact, almost surely.

Chapter 6

Mediated Mechanism Design for CPT Players

6.1 Introduction

In nearly every application of mechanism design, the decision-making entities are predominantly human beings faced with uncertainties. These uncertainties, for example, could arise from a combination of one or more factors from the following: (i) lack of information about the outcomes (e.g. oil lease auctions, kidney-exchange, insurance markets), (ii) each player having uncertainty about other players' behavior (e.g. voting behavior in elections, inclination to getting vaccinated in immunization programs), (iii) strategic interactions between the players (e.g. players could employ randomized strategies to hedge their market returns), (iv) randomness introduced by design (e.g. Tullock contests, where the probability of winning a prize depends on the amount of effort an agent puts into it). Naturally, to realize the mechanism designer's objectives, it is beneficial to consider as accurate and general models for human preference behavior under uncertainty as possible. Our goal here is to study mechanism design when players exhibit CPT preferences.

We are interested in situations where the *agents* participating in the system have private *types* (comprised of private information and preferences). The *system operator*¹⁸ is in a position to set the rules of communication and can control the implementation in the system. It aims to achieve certain goals, such as social welfare or revenue generation, without getting to directly observe the types of the players. Studying these systems when agents have CPT preferences requires modifications to the formal structures commonly encountered in classical mechanism design [58, 90, 93, 88, 83]. But before engaging in a systematic discussion of these issues, let us briefly describe our key result.

This starts with the observation that if the players are assumed to have CPT preferences instead of expected utility theory (EUT) preferences, then the revelation principle [92], one of the fundamental principles in mechanism design, does not hold anymore. A related observation was made in [71], where the authors show that in a second-price sealed-bid auction the

revelation principle holds in general if and only if the players have EUT preferences. Chew [35] provides an example to show that the revelation principle fails in a second-price sealed-bid auction when the players have preferences given by implicit weighted utility theory [40, 34].

The classical mechanism design framework is comprised of a fixed number of players, an allocation set, a set of types for each player, and a signal set for each player. (In this chapter, we will be concerned with the setting where all these sets are assumed to be finite.) The system operator commits to an allocation function, i.e. a function from the signal profile of the players to an allocation (see (6.2.11) for the formal definition).

The mechanism operates as follows:

1. Each player sends a signal strategically to the system operator based on its type (which is private knowledge to the player).
2. The system operator implements the allocation based on the signals from all the players in accordance to the allocation function that it committed to.

If we assume a prior over the types of the players which is common knowledge to all the agents and the system operator, and we assume that the signal sets of all the players, the allocation set and the allocation mapping are also common knowledge, then this constitutes a Bayesian game and one studies the outcome of such a game through its Bayes-Nash equilibria (see (6.2.15) for the formal definition). The revelation principle states that for the question of implementability of social choice functions (see (6.2.2) and (6.2.16) for formal definitions of social choice functions and their implementability), it is enough to assume the signal set to be the same as the type set for each player and confine attention to the equilibrium in which each player reports her type truthfully.

We propose a modification to the above framework that we call a *mediated mechanism*. We introduce a new stage where the system operator acts like a mediator and sends each player a private message sampled from a certain joint distribution on the set of message profiles. The allocation chosen by the system operator can now depend on both the message profile and the signal profile. Further, we explicitly allow the choice of the allocation to be randomized, which turns out to have no advantage in the classical mechanism design framework but can lead to benefits with CPT agents.

A mediated mechanism is therefore comprised of a fixed number of players, an allocation set, a set of types for each player, a message set for each player, and a signal set for each player, all of which are generally assumed to be finite sets. The system operator commits to a mediator distribution, which is a probability distribution on the set of message profiles. It also commits to a mediated allocation function, which maps each pair of signal profile and message profile to a probability distribution on allocations (see (6.4.2) for the formal definition).

The mechanism operates as follows:

1. The system operator samples a message profile from the declared mediator distribution and sends the individual messages to each player privately.

2. Each player receives her mediator message and, based on this message and her privately known type, sends a signal strategically to the system operator.
3. Based on the signals collected from all the players and the sampled message profile, the system operator samples the allocation in accordance to the probability distribution on allocations resulting from the mediated allocation function that it committed to.

Similarly to the previous setting we assume a prior over the types of the players that is common knowledge to all the agents and the system operator. We also assume that the message sets and the signal sets of all the players, the mediator distribution, the allocation set, and the allocation-outcome mapping are common knowledge. This along with the mechanism operation stated above constitutes a Bayesian game and we study the outcome of such a game through its Bayes-Nash equilibria (see (6.4.7) for the formal definition). With this modified framework, we recover a form of the revelation principle which states that it is enough to assume the signal set to be the same as the type set for each player and confine our attention to the equilibrium in which each player reports her true type irrespective of the private message she receives from the mediator. (See statement (i) of Theorem 6.4.1.)

As the mediator message sets could be arbitrary, it might seem that the problem of designing the signal sets has been transformed into the problem of designing the message sets. Although this is true, notice that the revelation principle allows us to restrict our attention to truthful strategies for each player, which have a simple form, thus resolving the difficult task of finding all the Bayes-Nash equilibria of the resulting game. Further, the fact that truthful reporting does not depend on the private message received by a player makes it a practical and natural strategy for the players.

We now resume our discussion of the different aspects involved in the study of mechanism design when agents have CPT preferences. The majority of the mechanism design literature has been restricted to EUT modeling of individual decision-making under uncertainty. Indeed, EUT has a nice normative interpretation and provides a useful and insightful first-order approximation (see, for example, [120]). However, systematic deviations from the predictions of EUT have been observed in several empirical studies involving human decision-makers [3, 48, 68] (see [124] for an excellent survey). With the advent of e-commerce activities and the ever-growing online marketplaces such as Amazon, eBay, and Uber, where the participating agents are largely human beings, who exhibit behavior that is highly susceptible to these deviations from EUT, it has become crucial to account for such behavioral deviations in the modeling of these systems. (For example, the paper [103] discusses the phenomenon of premium prices showing up in online marketplaces such as eBay to differentiate among sellers based on their reputation and buyers' perceived risks.)

A typical environment in the traditional mechanism design setup consists of a set of players that have private information about their types and an allocation set listing the possible alternatives from which the system operator chooses one that is best suited given the players' types. As mentioned earlier, we assume that the system operator controls the implementation and the players do not have separate decision domains. (Recall that by

private decision domains we mean possible actions for the player that directly affect the outcomes.) This is typical in several online marketplaces. For example, in online advertising platforms such as Google Ads, the platform has complete control over where to place which ads. Note that although the agents can affect the implementation of the system through their bids, these signals fall under the communication protocol set by the system, leaving the ultimate implementation in the hands of the system operator. In online matching markets such as eBay and Uber, the platform matches the buyers to sellers as in eBay, or riders to drivers as in Uber.¹⁹

Even if the system operator has complete control over the implementation, it wants the implementation to depend on the types of the agents. However, it does not have access to these types, and hence needs to design a mechanism to achieve this goal. Thus the system implementation indirectly depends on the choices of the participating agents. Note that the e-commerce applications mentioned above—Amazon, eBay, and Uber—fit well in this setup. Indeed, these are instances of a delivery system, an auction house, and a clearinghouse, which have been topics of interest for several years in mechanism design. However, the nature of these applications, and the presence of vast data corresponding to several repeated short-lived interactions of the system with any given user, makes it feasible to incorporate the behavioral features displayed by the users.

It has been a convention to assume that the outcome set for each player is identical to the allocation set, and hence the type for each player is assumed to capture her preferences over the allocation set (see, for example, [129]). However, in principle, the outcome set for any player need not be the same as the allocation set. Indeed the allocation set is a list of the alternatives available to the system operator to implement, whereas the outcome set consists of the outcomes realized by the players, and these can be quite different. For example, in the case of Amazon, the allocation set consists of alternative resource allocations to fulfill the delivery of purchased products, whereas the outcome set of a buyer consists of features such as time of delivery, place of delivery, etc. It makes sense to consider the preferences of a player over her outcome set, and any consideration of her preferences over the allocation set should be thought of as a pullback or a precomposition of her preferences over the outcome set with respect to the (possibly random) function that maps allocations to outcomes for this player.

We allow the above mapping from allocations to outcomes for any player to be randomized. Indeed, more often than not, the system operator does not have complete control over the outcomes of the players due to intrinsic uncertainties present in the system. For example, fixed resource allocations by Amazon can lead to uncertainty in the delivery times, possibly due to factors not part of the system model. In the case of Uber, upon matching the riders with the drivers in a certain way and choosing their corresponding routes, the arrival times and the riding experience of the users remain uncertain. In an auction setting such as eBay, if we consider the outcome set for any player to indicate if she receives the item or not, then the mapping from allocation to outcomes is deterministic. However, if we model the outcome set to indicate whether the player is satisfied with the item she receives, then we have to allow the mapping to be randomized.

Furthermore, the system operator might not be able to observe the outcome realization, for example the ride experience of a passenger. It can only try to learn this in hindsight through customer feedback. Besides, the outcome set for any single player is typically small as compared to the allocation set and the product of the outcome spaces of all the participating agents. Thus, treating each player's outcome set separately would enable us to focus on the preference behavior of an individual player and have better models for this player's preferences.

The (random) mapping from allocations to outcomes for any player induces a lottery L on the outcome set of this player for each allocation. EUT satisfies the linearity property which states that $U(\alpha L_1 + (1 - \alpha)L_2) = \alpha U(L_1) + (1 - \alpha)U(L_2)$, where $0 \leq \alpha \leq 1$, L_1, L_2 are two lotteries, and $U(\cdot)$ denotes the expected utility of the lottery within the parentheses. This property of EUT allows us to model the type of a player by considering her utility values for each allocation. For any lottery L over her outcomes that is induced by a lottery over the allocations μ , we can evaluate her utility $U(L)$ by taking the expectation over her utility values of the allocations with respect to the distribution μ . CPT on the other hand does not satisfy this linearity property (see, for example, [127]), and hence it is important that we consider the general model with separate outcome sets.

We formalize this general setup and provide preliminary background on CPT preferences in Section 6.2. Then, we consider the traditional mechanism design framework where each player knows her (private) type and strategically sends a signal to the system operator. The system operator collects these signals and implements a lottery over the allocation set.

We define a social choice function as a function mapping each type profile into a lottery over the product of the outcome sets for each player (see (6.2.2) for the formal definition). As an intermediate step, we consider an allocation choice function (i.e. a function that maps type profiles into lotteries over the allocation set, see (6.2.4)). Each allocation choice function uniquely defines a social choice function through the allocation-outcome mapping (see (6.2.5)), which we think of as a mapping from allocations to probability distributions on the product of the outcome sets of the agents. Note that there can be multiple allocation choice functions that give rise to the same social choice function. We define the notion of implementability for an allocation choice function in Bayes-Nash equilibrium (see (6.2.16)). We say that a social choice function is implementable in Bayes-Nash equilibrium if there exists an allocation choice function that is implementable in Bayes-Nash equilibrium and induces this social choice function.

We similarly define the notions of implementability in dominant equilibrium. Here, we identify an additional notion of implementability that we call implementable in belief-dominant equilibrium. Roughly speaking, a dominant strategy is a best response to all the strategy profiles of the opponents (see (6.2.18)), and a belief-dominant strategy is a best response to all the beliefs over the strategy profiles of the opponents (see (6.2.20)). Under EUT, the notion of a dominant strategy is equivalent to that of a belief-dominant strategy. However, this is not true in general when the agents have CPT preferences, thus making it necessary to distinguish between these two notions of equilibrium.

In Section 6.3, we define the notions of direct mechanism (see (6.3.1)) and truthful im-

plementation (see (6.3.2)). We then give an example that highlights the shortcoming of restricting oneself to direct mechanisms when the players have CPT preferences, as opposed to EUT preferences. In particular, we consider a 2-player setting where the players have CPT preferences that are not EUT preferences. Example 6.3.1 gives an allocation choice function for which the revelation principle does not hold for implementation in Bayes-Nash equilibrium. We then introduce the framework of mediated mechanism design in Section 6.4. We define the corresponding notions of Bayes-Nash equilibrium (see (6.4.7)), dominant equilibrium (see (6.4.11)), and belief-dominant equilibrium (see (6.4.12)) for mediated mechanisms. In Theorem 6.4.1, we recover the revelation principle under certain settings (see Table 6.1).

6.2 Mechanism Design Framework

Preliminaries

Let $[n] := \{1, 2, \dots, n\}$ be the set of *players* participating in the system. Let A denote the set of *allocations* for this system. We assume unless stated otherwise that the set of allocations is finite, say $A := \{\alpha^1, \dots, \alpha^l\}$. For example, in the sale of a single item (or multiple items), it could represent the allocation of the item(s) to the different individuals. In a routing system, such as traffic routing or internet packet routing, it could represent the different routing alternatives. More generally, in a resource allocation setting it could represent the assignment of resources to the participating agents (with their corresponding payments) that respect the system (and budget) constraints. In a voting scenario, it could represent the winning candidate. Thus, we imagine the allocations $\alpha \in A$ as being the various alternatives available to the system operator to implement.

Traditionally, each player is assumed to have a value for each of the allocations, and this defines the type of this player. It describes the preferences of a player over the allocations, and further, by assuming EUT behavior, we get her preferences over the lotteries over these allocations. Here, instead, we assume that for each player $i \in [n]$, we have a finite set of *outcomes* $\Gamma_i := \{\gamma_i^1, \dots, \gamma_i^{k_i}\}$, and player i 's *type* is defined by her CPT preferences over the lotteries on this set Γ_i . We imagine the set Γ_i to capture the outcome features that are relevant to player i . Thus the outcome set Γ_i allows us to separate out the features that affect player i from the underlying allocations that give rise to these outcomes. We capture this relation between the allocation set and the outcome sets through a mapping $\zeta : A \rightarrow \Delta(\Gamma)$ that we call the *allocation-outcome mapping*, where $\Gamma := \prod_i \Gamma_i$. Let $\zeta_i : A \rightarrow \Delta(\Gamma_i)$ denote allocation-outcome mapping for player i given by the marginal of ζ on the set Γ_i .

From a behavioral point of view it is natural to model a player's preferences on the outcome set Γ_i rather than the allocation set A . Then why is it that the sets Γ_i and the mapping ζ are usually missing from the mechanism design framework prevalent in the literature? At the end of this section after setting up the relevant notation, we will show that under EUT, from the point of view of the typical goals of the mechanism designer, it is enough to consider a transformation of the system where $\Gamma_i = A$, for all i , and the

allocation-outcome mappings are *trivial*, namely, $\zeta_i(\alpha) = \alpha$, for all $\alpha \in A, i \in [n]$ (this is shown formally in Appendix 6.C). We will also show that this does not hold in general when the players do not have EUT preferences, and in particular when they have CPT preferences.

We model the preference behavior of the players using cumulative prospect theory. Suppose Γ_i is the outcome set for player i , who is associated with a *value function* $v_i : \Gamma_i \rightarrow \mathbb{R}$ and two *probability weighting functions* $w_i^\pm : [0, 1] \rightarrow [0, 1]$. The value function v_i partitions the set of outcomes Γ_i into two parts: *gains* and *losses*; an outcome $\gamma_i \in \Gamma_i$ is said to be a gain if $v_i(\gamma_i) \geq 0$, and a loss otherwise. The probability weighting functions w_i^+ and w_i^- will be used for gains and losses, respectively. The probability weighting functions w_i^\pm are assumed to satisfy the following: (i) they are strictly increasing, (ii) $w_i^\pm(0) = 0$ and $w_i^\pm(1) = 1$. We say that (v_i, w_i^\pm) are the CPT features of player i .

Mechanism design framework

For each i , let Θ_i denote the set from which the permissible types for player i are drawn. Corresponding to any type θ_i for player i , let $v_i : \Gamma_i \rightarrow \mathbb{R}$ be her *value function*, and $w_i^\pm : [0, 1] \rightarrow [0, 1]$ be her *probability weighting functions*. Let $V_i^{\theta_i}(L_i)$ denote the CPT value of the lottery $L_i \in \Delta(\Gamma_i)$ for player i having type θ_i . Thus, the type θ_i completely determines the preferences of player i over lotteries on her outcome set Γ_i .²⁰ We will assume that the sets Θ_i are finite for all i .

Let $\theta := (\theta_1, \dots, \theta_n)$ denote the profile of types of the players, and let $\Theta := \prod_i \Theta_i$. We assume that each player knows her type but cannot observe the types of her opponents.

Let the set of players $[n]$, their corresponding type sets $\Theta_i, i \in [n]$, the allocation set A , and the outcome spaces $\Gamma_i, i \in [n]$, together with the mapping ζ form an *environment*, denoted by

$$\mathcal{E} := ([n], (\Theta_i)_{i \in [n]}, A, (\Gamma_i)_{i \in [n]}, \zeta). \quad (6.2.1)$$

A *social choice function*

$$g : \Theta \rightarrow \Delta(\Gamma) \quad (6.2.2)$$

determines a lottery over the product of the outcome sets of the individual players given the type profile θ of all the players. The *outcome choice function* for player i corresponding to the social choice function g is

$$g_i : \Theta \rightarrow \Delta(\Gamma_i), \quad (6.2.3)$$

given by the restriction of g to the set Γ_i , and represents the lottery faced by player i given the type profile θ of all the players. We will treat the social choice function g as the goal of the mechanism designer, i.e, the goal is to design a mechanism to implement a social choice function g without having knowledge of the true types of the players.

Let an *allocation choice function*

$$f : \Theta \rightarrow \Delta(A) \quad (6.2.4)$$

represent the choice of the allocation to be implemented by the system operator given a type profile $\theta \in \Theta$. Note that $f(\theta)$ is a probability distribution over the allocations A .

Thus we allow the system operator to implement a randomized allocation. A *deterministic allocation choice function* maps each type profile to a unique allocation. Since the mapping ζ is fixed and a part of the environment description, the allocation choice function f effectively captures the goal of a mechanism designer. More precisely, let $\mathcal{F}(g)$ denote the set of all allocation choice functions that induce the social choice function g , i.e. for all $\theta \in \Theta$, $g(\theta)$ is the mixture probability distribution of the probability distributions $(\zeta(\alpha), \alpha \in A)$ with weights $f(\alpha|\theta)$. We note that the set $\mathcal{F}(g)$ is non-empty if and only if

$$g(\theta) \in \text{co}\{\zeta(\alpha) : \alpha \in A\},$$

for all $\theta \in \Theta$. We wish to design a mechanism that would implement an allocation choice function in $\mathcal{F}(g)$. Thus a social choice function is implementable if and only if we can implement an allocation choice function f that satisfies

$$g(\gamma|\theta) = \sum_{\alpha \in A} f(\alpha|\theta)\zeta(\gamma|\alpha), \quad (6.2.5)$$

for all $\gamma \in \Gamma, \theta \in \Theta$. This raises the main question in mechanism design, namely whether we can design a game that results in the implementation of some given allocation choice function f under certain rationality conditions on the players even when the system operator cannot observe the players' types.

First, let us look at the the relationship between lotteries on allocations and lotteries on the outcome set of a given player. Any lottery $\mu \in \Delta(A)$ induces a lottery $L_i(\mu) \in \Delta(\Gamma_i)$ given by

$$L_i(\gamma_i|\mu) := \sum_{\alpha \in A} \mu(\alpha)\zeta_i(\gamma_i|\alpha). \quad (6.2.6)$$

Given that player i has type θ_i , we know that the CPT value of lottery $L_i(\mu)$ is $V_i^{\theta_i}(L_i(\mu))$. This induces a value for player i with type θ_i on the lottery μ denoted by

$$W_i^{\theta_i}(\mu) := V_i^{\theta_i}(L_i(\mu)). \quad (6.2.7)$$

This defines a *utility function* $W_i^{\theta_i} : \Delta(A) \rightarrow \mathbb{R}$ that gives the preference relation over the lotteries $\mu \in \Delta(A)$ for a player i having type θ_i . Let

$$u_i^{\theta_i}(\alpha) := V_i^{\theta_i}(\zeta_i(\alpha)) = W_i^{\theta_i}(\alpha) \quad (6.2.8)$$

be the CPT value of the lottery for player i corresponding to allocation α .²¹ If player i has EUT preferences, then we have that

$$W_i^{\theta_i}(\mu) = \sum_{\alpha \in A} \mu(\alpha)u_i^{\theta_i}(\alpha). \quad (6.2.9)$$

We now consider a *mechanism*

$$\mathcal{M}_0 := ((\Psi_i)_{i \in [n]}, h_0), \quad (6.2.10)$$

consisting of a collection of finite *signal sets* Ψ_i , one for each player i , and an *allocation function*

$$h_0 : \Psi \rightarrow \Delta(A), \quad (6.2.11)$$

where $\Psi := \prod_{i \in [n]} \Psi_i$. Note that the allocation function is allowed to be randomized. Let ψ_i denote a typical element of Ψ_i , and $\psi := (\psi_i)_{i \in [n]}$ denote a typical element of Ψ , called a *signal profile*.

It is straightforward to incorporate the feature that the outcome sets Γ_i might be different from the allocation set A , and the corresponding allocation-outcome mapping ζ , so as to extend the definition of a Bayes-Nash equilibrium strategy profile for the mechanism \mathcal{M}_0 and the implementability of an allocation choice function f in Bayes-Nash equilibrium. To do this, assume that the types of the individual players are drawn according to a prior distribution $F \in \Delta(\Theta)$ and that this distribution is common knowledge among the agents and the system operator. Let $F_i \in \Delta(\Theta_i)$ denote the marginal of F on Θ_i . Suppose player i has type θ_i . Then the belief of player i about the types of other players is given by the conditional distribution

$$F_{-i}(\theta_{-i}|\theta_i) := \frac{F(\theta_i, \theta_{-i})}{F_i(\theta_i)}, \text{ for all } \theta_{-i} \in \Theta_{-i}, \theta_i \in \text{supp } F_i,$$

where $\theta_{-i} := (\theta_j)_{j \neq i}$ is the profile of types of all players other than player i , $\Theta_{-i} := \prod_{j \neq i} \Theta_j$.

Recall that ψ_i denotes a typical element of Ψ_i , and $\psi := (\psi_i)_{i \in [n]}$ denotes a typical element of Ψ . Let $\Psi_{-i} := \prod_{j \neq i} \Psi_j$ with a typical element denoted by ψ_{-i} . Let

$$\sigma_i : \Theta_i \rightarrow \Delta(\Psi_i) \quad (6.2.12)$$

be a *strategy* for player i , and let $\sigma := (\sigma_1, \sigma_2, \dots, \sigma_n)$ denote a strategy profile. Let $\sigma_{-i} := (\sigma_j)_{j \neq i}$ denote the strategy profile of all players other than player i . For any type θ_i (such that $F_i(\theta_i) > 0$) and signal ψ_i , consider the probability distribution $\mu_i(\theta_i, \psi_i; \mathcal{M}_0, F, \sigma_{-i}) \in \Delta(A)$ given by

$$\mu_i(\alpha|\theta_i, \psi_i; \mathcal{M}_0, F, \sigma_{-i}) := \sum_{\theta_{-i} \in \Theta_{-i}} F_{-i}(\theta_{-i}|\theta_i) \sum_{\psi_{-i} \in \Psi_{-i}} \prod_{j \neq i} \sigma_j(\psi_j|\theta_j) h_0(\alpha|\psi), \quad (6.2.13)$$

for all $\alpha \in A$. Suppose player i has type θ_i (such that $F_i(\theta_i) > 0$), and she chooses to signal ψ_i . Then, by Bayes' rule, the lottery faced by player i is given by

$$L_i(\mu_i(\theta_i, \psi_i; \mathcal{M}_0, F, \sigma_{-i})).$$

This comes from the assumption that player i knows her type θ_i , the common prior F , the strategies $\sigma_j, j \neq i$ of her opponents, and the mapping ζ_i . Given that player i has type θ_i , her CPT value for the lottery $L_i(\mu_i(\theta_i, \psi_i; \mathcal{M}_0, F, \sigma_{-i}))$ is given by

$$W_i^{\theta_i}(\mu_i(\theta_i, \psi_i; \mathcal{M}_0, F, \sigma_{-i})) = V_i^{\theta_i}(L_i(\mu_i(\theta_i, \psi_i; \mathcal{M}_0, F, \sigma_{-i})),$$

where we recall that $W_i^{\theta_i}(\mu)$ is the CPT value of player i with type θ for the lottery $L_i(\mu) \in \Delta(\Gamma_i)$ induced by the distribution $\mu \in \Delta(A)$. Let the *best response strategy set* $BR_i(\sigma_{-i})$ for player i to a strategy profile σ_{-i} of her opponents consist of all strategies $\sigma_i^* : \Theta_i \rightarrow \Delta(\Psi_i)$ such that

$$W_i^{\theta_i}(\mu_i(\theta_i, \psi_i; \mathcal{M}_0, F, \sigma_{-i})) \geq W_i^{\theta_i}(\mu_i(\theta_i, \psi'_i; \mathcal{M}_0, F, \sigma_{-i})), \quad (6.2.14)$$

for all $\theta_i \in \text{supp } F_i$, $\psi_i \in \text{supp } \sigma_i^*(\theta_i)$, $\psi'_i \in \Psi_i$. The rationale behind this definition is that a player's best response strategy σ^* should not assign positive probability to any suboptimal signal ψ_i given her type θ_i .

A strategy profile σ^* is said to be an *F-Bayes-Nash equilibrium* for the environment \mathcal{E} and common prior F with respect to the mechanism \mathcal{M}_0 if, for each player i , we have

$$\sigma_i^* \in BR_i(\sigma_{-i}^*). \quad (6.2.15)$$

We will refer to σ^* simply as a Bayes-Nash equilibrium when the respective environment \mathcal{E} , the common prior F , and mechanism \mathcal{M}_0 are clear from the context.

We say that the allocation choice function f is *implementable in F-Bayes-Nash equilibrium* by a mechanism if there exists a mechanism \mathcal{M}_0 and an F-Bayes-Nash equilibrium σ such that f is the induced distribution, i.e. for all $\theta_i \in \text{supp } F_i$, $\alpha \in A$, we have

$$f(\alpha|\theta) = \sum_{\psi \in \Psi} \left(\prod_{i \in [n]} \sigma_i(\psi_i|\theta_i) \right) h_0(\alpha|\psi). \quad (6.2.16)$$

An alternative notion is that of an allocation choice function f being *implementable in dominant equilibrium*. The traditional notion states that a strategy σ_i is a *dominant strategy* for player i if the signals in the support of $\sigma_i(\theta_i)$ are optimal given player i 's type θ_i and any signal profile ψ_{-i} of the opponents. More precisely, if we let

$$\mu_i(\theta_i, \psi_i; \mathcal{M}_0, \psi_{-i}) := h_0(\psi_i, \psi_{-i}), \quad (6.2.17)$$

then σ_i^* is a dominant strategy if, for all $\theta_i \in \Theta_i$, $\psi_i \in \text{supp } \sigma_i^*(\theta_i)$, $\psi'_i \in \Psi_i$, and $\psi_{-i} \in \Psi_{-i}$, we have

$$W_i^{\theta_i}(\mu_i(\theta_i, \psi_i; \mathcal{M}_0, \psi_{-i})) \geq W_i^{\theta_i}(\mu_i(\theta_i, \psi'_i; \mathcal{M}_0, \psi_{-i})). \quad (6.2.18)$$

Thus, if player i employs a dominant strategy, then regardless of the signal profile of the opponents she always signals a best response given her type. A dominant equilibrium is one in which each player plays a dominant strategy. We say that an allocation choice function f is implementable in dominant equilibrium if there exists a mechanism \mathcal{M}_0 and a strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ where each σ_i^* is a dominant strategy (equivalently, σ^* is a dominant equilibrium) such that (6.2.16) holds for all $\theta_i \in \Theta$, $\alpha \in A$.

Under EUT, if a signal ψ_i is a best response of player i for all of the opponents' signal profiles, then it is also a best response for any belief $G_{-i} \in \Delta(\Psi_{-i})$ of player i over her

opponents' signal profiles. However, under CPT, this need not hold. (See Example 6.2.1.) This observation leads us to the following stricter notion of dominant strategies under CPT. We call a strategy σ_i a *belief-dominant strategy* for player i if the signals in the support of $\sigma_i(\theta_i)$ are optimal given player i 's type θ_i and any belief $G_{-i} \in \Delta(\Psi_{-i})$ she has on the signal profile of her opponents. If we let

$$\mu_i(\theta_i, \psi_i; \mathcal{M}_0, G_{-i}) := \sum_{\psi_{-i}} G_{-i}(\psi_{-i}) h_0(\psi_i, \psi_{-i}), \quad (6.2.19)$$

then ψ_i^* is a belief-dominant strategy for player i if, for all $\theta_i \in \Theta_i$, $\psi_i \in \text{supp } \sigma_i^*(\theta_i)$, $\psi_i' \in \Psi_i$, and $G_{-i} \in \Delta(\Psi_{-i})$, we have

$$W_i^{\theta_i}(\mu_i(\theta_i, \psi_i; \mathcal{M}_0, G_{-i})) \geq W_i^{\theta_i}(\mu_i(\theta_i, \psi_i'; \mathcal{M}_0, G_{-i})). \quad (6.2.20)$$

It is straightforward to check that under EUT a strategy is dominant if and only if it is belief-dominant. A *belief-dominant equilibrium* is one in which every player plays a belief-dominant strategy. We say that an allocation choice function f is implementable in belief-dominant equilibrium if there exists a mechanism \mathcal{M}_0 and a strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ where each σ_i^* is a belief-dominant strategy (equivalently, σ^* is a belief-dominant equilibrium) such that (6.2.16) holds for all $\theta_i \in \Theta, \alpha \in A$.

Note that if σ^* is a belief-dominant strategy profile, and thus a belief-dominant equilibrium, then it is a dominant strategy profile, i.e. a dominant equilibrium, and also an F -Bayes-Nash equilibrium with respect to any prior distribution F on type profiles.

Example 6.2.1. Let $n = 2$. Let $\Theta_1 = \Theta_2 = \{\text{UP}, \text{DN}\}$. Let $A = \{a, b, c\}$, $\Gamma_1 = \{\text{I}, \text{II}, \text{III}, \text{IV}, \text{V}\}$, and $\Gamma_2 = A$. Let the allocation-outcome mapping be given by the product distribution of the marginals ζ_1 and ζ_2 , given by $\zeta_1(a) = \{(1/2, \text{I}); (1/2, \text{V})\}$, $\zeta_1(b) = \{(1/2, \text{II}); (1/2, \text{IV})\}$, $\zeta_1(c) = \{(1, \text{III})\}$, and $\zeta_2(\alpha) = \alpha, \forall \alpha \in A$. Let the probability weighting functions for gains for the two players be given by

$$w_1^+(p) = \exp\{-(-\ln p)^{0.5}\}, w_2^+(p) = p,$$

for $p \in [0, 1]$. In this example, we consider only lotteries with outcomes in the gains domain, and hence an arbitrary probability weighting function for the losses can be assumed. Here, for player 1's probability weighting function, we employ the form suggested by Prelec [113] (see Figure 6.1). Note that player 2 has EUT preferences. Let the value functions v_1 and v_2 be given by

v_1	I	II	III	IV	V
UP	$2x$	$x + 1$	1.99	1	0
DN	0	0	1	0	0

v_2	a	b	c
UP	1	0	2
DN	0	1	2

where $x := 1/w_1^+(0.5) = 2.2992$. Note that $2x = 4.5984$ and $x + 1 = 3.2992$. We have,

$$\begin{aligned} V_1^{\text{UP}}(L_1(a)) &= 2xw_1^+(0.5) = 2, \\ V_1^{\text{UP}}(L_1(b)) &= 1 + xw_1^+(0.5) = 2, \end{aligned}$$




Figure 6.1: The solid curve shows the probability weighting function w_1^+ for player 1 from Example 6.2.1 and Example 6.4.2. The dotted curve shows the probability weighting function w_2^+ for player 2 in Example 6.2.1 and Example 6.4.2, which is the linear function corresponding to EUT preferences.

and,

$$\begin{aligned} V_1^{\text{UP}}(0.5L_1(a) + 0.5L_1(b)) &= w_1^+(0.75) + xw_1^+(0.5) + (x-1)w_1^+(0.25) \\ &= 1.9851. \end{aligned}$$

(Here, we have $w_1^+(0.25) = 0.3081$, $w_1^+(0.5) = 0.4349$, and $w_1^+(0.75) = 0.5849$.) Consider the mechanism $\mathcal{M} = ((\Psi_i)_{i \in [n]}, h_0)$, where $\Psi_1 = \Psi_2 = \{\text{UP}, \text{DN}\}$, and h_0 is given by

$$h_0(\text{UP}, \text{UP}) = a, h_0(\text{UP}, \text{DN}) = b, h_0(\text{DN}, \text{UP}) = c, h_0(\text{DN}, \text{DN}) = c.$$

Consider the strategies

$$\sigma_i(\text{UP}) = \text{UP}, \text{ and } \sigma_i(\text{DN}) = \text{DN},$$

for both the players i . It is easy to see that these strategies are dominant for both the players. However, if player 1 has type UP and believes that there is an equal chance of player 2 reporting her strategy to be UP and DN, then player 1's best response is to report DN. Thus, σ_1 is not a belief-dominant strategy for player 1. □

We will now look at the remark made earlier about the absence of the distinction between the allocation set and the outcome sets in classical mechanism design, and why it is important to consider this distinction under CPT. In Appendix 6.C, we show that under EUT it suffices to consider the scenario where the outcome set of each player is the same as the allocation

set by the simple expedient of interpreting each type $\theta_i \in \Theta_i$ in terms of the utility function on allocations that it defines via (6.2.8).

While equation (6.2.9) holds under EUT, under CPT in general it does not hold, and in general the utility function $W_i^{\theta_i}$ is not completely determined by the values $u_i^{\theta_i}(\alpha), \forall \alpha \in A$. Thus, we can either characterize the type of a player by her utility function $W_i^{\theta_i}$ or by her CPT features which, combined with the mapping ζ_i , together define the utility function $W_i^{\theta_i}$. In any given setting, it is more natural to put behavioral assumptions on the CPT features (v_i, w_i^\pm) than on the utility function $W_i^{\theta_i}$.²² Hence, we include the sets Γ_i and the mappings ζ_i , for all i , in our system model under CPT.

6.3 The Revelation Principle

A mechanism $\mathcal{M}_0 = ((\Psi_i)_{i \in [n]}, h_0)$ is called a *direct mechanism* if $\Psi_i = \Theta_i$, for all i . Let $\mathcal{M}_0^d := ((\Theta_i)_{i \in [n]}, h_0^d)$ denote a direct mechanism, where

$$h_0^d : \Theta \rightarrow \Delta(A) \quad (6.3.1)$$

is the *direct allocation function*. Corresponding to a direct mechanism, let $\sigma_i^d : \Theta_i \rightarrow \Theta_i$ denote the truthful strategy for player i , given by

$$\sigma_i^d(\theta_i) = \theta_i, \quad (6.3.2)$$

for all $\theta_i \in \Theta_i$. An allocation choice function f is said to be *truthfully implementable* in F -Bayes-Nash equilibrium (resp. dominant equilibrium or belief-dominant equilibrium) if there exists a direct mechanism \mathcal{M}_0^d such that the truthful strategy profile σ^d is an F -Bayes-Nash equilibrium (resp. dominant equilibrium or belief-dominant equilibrium), and it induces f .

The *revelation principle*²³ says that if an allocation choice function is implementable in Bayes-Nash equilibrium (resp. dominant equilibrium or belief-dominant equilibrium) by a mechanism, then it is also truthfully implementable in Bayes-Nash equilibrium (resp. dominant equilibrium or belief-dominant equilibrium) by a direct mechanism. When the players are restricted to have EUT preferences and the outcome set of each player is assumed to be the same as the allocation set with the trivial allocation-outcome mapping, Myerson [93] proved that the revelation principle holds for both the versions - Bayes-Nash equilibrium and dominant equilibrium (and hence also for belief-dominant equilibrium, since dominant strategies are equivalent to belief-dominant strategies under EUT). It is easy to extend this result to the general setting where some of the individual outcome sets might differ from the allocation set, provided the players are restricted to have EUT preferences. Indeed, in Appendix 6.C it is proved that, under EUT, an allocation choice function f is implementable in F -Bayes-Nash (resp. dominant or belief-dominant) equilibrium by a mechanism \mathcal{M}_0 for the environment \mathcal{E} with the equilibrium strategy σ , if and only if, for the corresponding environment \mathcal{E}' (defined in Appendix 6.C), the corresponding allocation choice function f' is implementable in F' -Bayes-Nash (resp. dominant or belief-dominant) by the same

mechanism \mathcal{M}_0 with the corresponding equilibrium strategy σ' . We now observe that \mathcal{M}_0 is a direct mechanism for environment \mathcal{E} if and only if it is a direct mechanism for environment \mathcal{E}' . Also, σ_i is the truthful strategy with respect to the environment \mathcal{E} and a direct mechanism \mathcal{M}_0 , if and only if, the corresponding strategy σ'_i is the truthful strategy with respect to the environment \mathcal{E}' and the same direct mechanism \mathcal{M}_0 . These observations together give us the required revelation principle under EUT for the setting where the outcome sets of some of the players can differ from the allocation set.

The following example shows that the revelation principle need not hold when players have CPT preferences. We will consider implementability in Bayes-Nash equilibrium in this example.

Example 6.3.1. Let there be two players, i.e. $n = 2$. Let each player belong to one of the three types: Mildly Favorable (MF), Unfavorable (UF), and Super Favorable (SF), i.e. $\Theta_1 = \Theta_2 = \{\text{MF}, \text{UF}, \text{SF}\}$. Let the outcome sets for both the players be $\Gamma_1 = \Gamma_2 = \{\text{I}, \text{II}, \text{III}, \text{IV}, \text{V}\}$. Let the value functions v_1 and v_2 for both the players be as shown below.

	I	II	III	IV	V
MF	13.616	8.616	5.816	3.8	0
UF	-190	-100	-1K	-50	0
SF	0	0	1M	0	0

Observe that a player with type MF has mild gains for all the outcomes, a player with type UF has medium losses for all outcomes except outcome III, where she has a big loss, and a player of type SF has a huge gain for outcome III and zero gains otherwise.

Let the probability weighting functions for both the players, for all of their types, be given by the following piecewise linear functions:

$$w^+(p) = \begin{cases} (8/7)p, & \text{for } 0 \leq p < (7/32), \\ (1/4) + (2/3)(p - 7/32), & \text{for } (7/32) \leq p < 25/32, \\ (5/8) + (12/7)(p - 25/32), & \text{for } (25/32) \leq p < 1, \end{cases}$$

for gains, and

$$w^-(p) = \begin{cases} (3/2)p, & \text{for } 0 \leq p < (1/8), \\ (3/16) + (1/2)(p - 1/8), & \text{for } (1/8) \leq p < 3/4, \\ (1/2) + 2(p - 3/4), & \text{for } (3/4) \leq p < 1, \end{cases}$$

for losses. See Figure 6.1.

Let the prior distribution F be such that the types of the players are independently sampled with probabilities,

$$\mathbb{P}(\text{MF}) = 1/2, \mathbb{P}(\text{UF}) = 3/8, \mathbb{P}(\text{SF}) = 1/8. \tag{6.3.3}$$




Figure 6.1: Probability weighting functions for the players in Example 6.3.1.

Let $A = \{a, b, c\}$ be the allocation set, and let the allocation-outcome mapping be given by

$$\begin{aligned}\zeta(a) &= \{(1/2, (\text{I}, \text{I})); (1/2, (\text{V}, \text{V}))\}, \\ \zeta(b) &= \{(1/2, (\text{II}, \text{II})); (1/2, (\text{IV}, \text{IV}))\}, \\ \zeta(c) &= (\text{III}, \text{III}).\end{aligned}$$

Consider the allocation choice function f^* given by

$$\begin{aligned}f^*(\text{SF}, \theta_2) &= f^*(\theta_1, \text{SF}) = c, \quad \forall \theta_1 \in \Theta_1, \theta_2 \in \Theta_2, \\ f^*(\text{UF}, \theta_2) &= f^*(\theta_1, \text{UF}) = \{(1/2, a); (1/2, b)\}, \quad \forall \theta_1 \in \{\text{MF}, \text{UF}\}, \theta_2 \in \{\text{MF}, \text{UF}\}, \\ f^*(\text{MF}, \text{MF}) &= \{(1/2, a); (1/2, b)\}.\end{aligned}$$

We will now show that f^* is *not* truthfully implementable in F -Bayes-Nash equilibrium by a direct mechanism. However, if we do not restrict ourselves to direct mechanisms, then we will show that it is possible to implement f^* in F -Bayes-Nash equilibrium. We will then conclude that the revelation principle does not hold for Bayes-Nash implementability when the players have CPT preferences.

We observe that if either of the players is of type SF then under the allocation c the players with type SF get the maximum possible reward, i.e. $1M$. This motivates implementing allocation c if either of the players is of type SF. Now suppose none of the players has type SF. If player 1 is of type UF, then in Claim 6.3.2, we show that player 1's CPT value for the lottery $L_i(\mu)$ corresponding to a distribution $\mu \in \Delta(A)$ is maximized when

$$\mu = \{(1/2, a); (1/2, b); (0, c)\}. \quad (6.3.4)$$

Thus, if at least one of the players has type UF and none of the players have type SF, then the distribution in (6.3.4) gives the best CPT value for the players with type UF. This motivates the following definition: we will call an allocation choice function f *special* if it satisfies

$$f(\text{SF}, \theta_2) = f(\theta_1, \text{SF}) = \{(1, c)\}, \forall \theta_1 \in \Theta_1, \theta_2 \in \Theta_2, \quad (6.3.5)$$

and

$$f(\text{UF}, \theta_2) = f(\theta_1, \text{UF}) = \{(1/2, a); (1/2, b)\}, \forall \theta_1, \theta_2 \in \{\text{MF}, \text{UF}\}. \quad (6.3.6)$$

Note that f^* is special.

After proving Claim 6.3.2, we will show that it is impossible to truthfully implement any special allocation choice function in F -Bayes-Nash equilibrium by a direct mechanism. In particular, this would imply that f^* is not truthfully implementable by a direct mechanism. We will then give a mechanism \mathcal{M}_0 that implements f^* in F -Bayes-Nash equilibrium.

Claim 6.3.2. The CPT value $V_1^{\text{UF}}(L_1(\mu))$ is maximized when μ is given by (6.3.4).

Proof of Claim 6.3.2. Consider a lottery

$$\mu = \{(x, a); (y, b); (z, c)\},$$

where x, y, z are nonnegative numbers with $x + y + z = 1$. Then the outcome lottery for player 1 is

$$L_1(\mu) = \{(x/2, \text{I}); (y/2, \text{II}); (z, \text{III}); (y/2, \text{IV}); (x/2, \text{V})\}.$$

CPT satisfies *strict stochastic dominance* [30], i.e. shifting positive probability mass from an outcome to a strictly preferred outcome leads to a strictly preferred lottery. This implies that $z = 0$ in the optimal solution. Taking $z = 0$ and $y = 1 - x$, from (1.3.9), we have

$$\begin{aligned} E(x) &:= V_1^{\text{UF}}(\{(x/2, \text{I}); (1/2 - x/2, \text{II}); (0, \text{III}); (1/2 - x/2, \text{IV}); (x/2, \text{V})\}) \\ &= -90w^-(x/2) - 50w^-(1/2) - 50w^-(1 - x/2). \end{aligned}$$


We can verify that this function is maximized at $x = 1/2$. See Figure 6.2. □

Suppose we have a direct mechanism $\mathcal{M}_0^d = h_0^d$ that truthfully implements a special allocation choice function f . Then the allocation function h_0^d must be equal to the allocation choice function f . Since f satisfies (6.3.5) and (6.3.6), the only freedom left is in the choice of $f(\text{MF}, \text{MF})$. Let

$$h_0^d(\text{MF}, \text{MF}) = f(\text{MF}, \text{MF}) = \{(x', a); (y', b); (z', c)\},$$

where x', y', z' are nonnegative numbers with $x' + y' + z' = 1$. The lottery faced by a player of type MF signaling truthfully would then be

$$\begin{aligned} L_1(\mu_1(\text{MF}, \text{MF}; \mathcal{M}_0^d, F, \sigma_{-i}^d)) &= \{(3/32 + x'/4, \text{I}); (3/32 + y'/4, \text{II}); \\ &\quad (1/8 + z'/2, \text{III}); (3/32 + y'/4, \text{IV}); (3/32 + x'/4, \text{V})\}. \end{aligned}$$

Figure 6.2: Plot of expression $E(x)$ in Claim 6.3.2.

We obtain this by using the belief $F_{-1}(\cdot|\text{MF})$ of player 1 on the type of player 2 given by (6.3.3), the truthful strategy σ_2^d for player 2, and the allocation function h_0^d in (6.2.13).

Claim 6.3.3. For any nonnegative x', y', z' such that $x' + y' + z' = 1$, we have

$$V_1^{\text{MF}}(L_1(\mu_1(\text{MF}, \text{MF}; \mathcal{M}_0^d, F, \sigma_{-i}^d))) < 5.816.$$

Proof of Claim 6.3.3. We have

$$\begin{aligned} V_1^{\text{MF}} L_1(\mu_1(\text{MF}, \text{MF}; \mathcal{M}_0^d, F, \sigma_{-i}^d)) &= 3.8w^+(29/32 - x'/4) + 2.016w^+(18/32 + z'/4) \\ &\quad + 2.8w^+(14/32 - z'/4) + 5w^+(3/32 + x'/4). \end{aligned}$$

We observe that the expression,

$$E_1(z') := 2.016w^+(18/32 + z'/4) + 2.8w^+(14/32 - z'/4),$$

is maximized at $z' = 0$ with value $E_1(0) = 2.0743$. See Figure 6.3.

We can therefore set $z' = 0$, since this choice would also lead to the least constrained problem of maximizing the expression

$$E_2(x') := 3.8w^+(29/32 - x'/4) + 5w^+(3/32 + x'/4) + 2.0743,$$

which we can see is maximized at $x' = 0$ and $x' = 1$. At $z' = 0$, and either $x' = 0$ or $x' = 1$, we have $V_1^{\text{MF}} L_1(\mu_1(\text{MF}, \text{MF})) = 5.7993$. See Figure 6.4.

This establishes the claim. \square




Figure 6.3: Plot of expression $E_1(z')$ in Claim 6.3.3.




Figure 6.4: Plot of expression $E_2(x')$ in Claim 6.3.3.

Thus, player 1 will defect from the truthful strategy and report SF when her true type is MF, because if she does so the allocation c will be implemented by the system operator, which results in her outcome being III, hence giving her a value of 5.816. Hence, truthful strategies do not form a Bayes-Nash equilibrium under the allocation function h_0^d . And hence, any allocation choice function f that satisfies (6.3.5) and (6.3.6) is not truthfully implementable by a direct mechanism.

We will now show that the allocation choice function f^* is implementable in Bayes-Nash equilibrium. Consider the mechanism $\mathcal{M}_0 = ((\Psi_i)_i, h_0)$ with the signal sets for the players being $\Psi_1 = \Psi_2 = \{\text{MF}^a, \text{MF}^b, \text{UF}, \text{SF}\}$, and the allocation function h_0 given by

$$\begin{aligned} h_0(\text{SF}, \psi_2) &= h_0(\psi_1, \text{SF}) = c, \quad \forall \psi_1 \in \Psi_1, \psi_2 \in \Psi_2, \\ h_0(\text{UF}, \text{UF}) &= \{(1/2, a); (1/2, b)\}, \\ h_0(\text{UF}, \text{MF}^a) &= a, \\ h_0(\text{UF}, \text{MF}^b) &= b, \\ h_0(\text{MF}^a, \text{UF}) &= a, \\ h_0(\text{MF}^b, \text{UF}) &= b, \\ h_0(\text{MF}^a, \text{MF}^a) &= a, \\ h_0(\text{MF}^b, \text{MF}^b) &= b, \\ h_0(\text{MF}^a, \text{MF}^b) &= h_0(\text{MF}^b, \text{MF}^a) = \{(1/2, a); (1/2, b)\}. \end{aligned}$$

Now consider the strategies σ_1^* and σ_2^* given by

$$\begin{aligned} \sigma_i^*(\text{SF}) &= \text{SF}, \\ \sigma_i^*(\text{UF}) &= \text{UF}, \\ \sigma_i^*(\text{MF}) &= \{(1/2, \text{MF}^a); (1/2, \text{MF}^b)\}, \end{aligned} \tag{6.3.7}$$

for $i = 1, 2$.

One can check that the allocation function h_0 and the strategy profile σ^* induce the allocation choice function f^* defined above. We will now verify that σ^* is a Bayes-Nash equilibrium and thus conclude that f^* is implementable in Bayes-Nash equilibrium.

If player i has type SF then clearly SF is a best response signal for her. To see this, observe that amongst all the lotteries $L_i \in \Delta(\Gamma_i)$, $V_i^{\text{SF}}(L_i)$ is maximized when $L_i = \text{III}$ (this follows from the first order stochastic dominance property of CPT preferences). Since signaling SF produces the lottery III for player i , we get that it is her best response. If player i has type UF, then signaling UF dominates signaling SF. To see this, note that amongst all the lotteries $L_i \in \Delta(\Gamma_i)$, $V_i^{\text{UF}}(L_i)$ is minimized when $L_i = \text{III}$ (this follows from the first order stochastic dominance property of CPT preferences). Since signaling SF produces the lottery III for player i , we get that it is dominated by all other strategies, in particular, signaling UF. As for comparing with signaling MF^a or MF^b , if she signals UF then she will

face the lottery

$$L_i(\mu_i(\text{UF}, \text{UF}; \mathcal{M}_0, F, \sigma_{-i}^*)) = \{(7/32, \text{I}); (7/32, \text{II}); (1/8, \text{III}); (7/32, \text{IV}); (7/32, \text{V})\}.$$

If she signals MF^a , then she will face the lottery

$$L_i(\mu_i(\text{UF}, \text{MF}^a; \mathcal{M}_0, F, \sigma_{-i}^*)) = \{(3/8, \text{I}); (1/16, \text{II}); (1/8, \text{III}); (1/16, \text{IV}); (3/8, \text{V})\}.$$

If she signals MF^b , then she will face the lottery

$$L_i(\mu_i(\text{UF}, \text{MF}^b; \mathcal{M}_0, F, \sigma_{-i}^*)) = \{(1/16, \text{I}); (3/8, \text{II}); (1/8, \text{III}); (3/8, \text{IV}); (1/16, \text{V})\}.$$

The CPT values in each of these scenarios are as follows:

$$\begin{aligned} & V_i^{\text{UF}}(L_i(\mu_i(\text{UF}, \text{UF}; \mathcal{M}_0, F, \sigma_{-i}^*))) \\ &= -50w^-(25/32) - 50w^-(18/32) - 90w^-(11/32) - 810w^-(4/32) \\ &= -227.0312, \end{aligned}$$

$$\begin{aligned} & V_i^{\text{UF}}(L_i(\mu_i(\text{UF}, \text{MF}^a; \mathcal{M}_0, F, \sigma_{-i}^*))) \\ &= -50w^-(20/32) - 50w^-(18/32) - 90w^-(16/32) - 810w^-(4/32) \\ &= -227.8125, \end{aligned}$$

and,

$$\begin{aligned} & V_i^{\text{UF}}(L_i(\mu_i(\text{UF}, \text{MF}^b; \mathcal{M}_0, F, \sigma_{-i}^*))) \\ &= -50w^-(30/32) - 50w^-(18/32) - 90w^-(6/32) - 810w^-(4/32) \\ &= -235.6250. \end{aligned}$$

Thus, signaling UF is the best response of a player with type UF.

Finally, let player i have type MF. Depending on what she signals, we have the following lotteries:

$$\begin{aligned} L_i(\mu_i(\text{MF}, \text{MF}^a; \mathcal{M}_0, F, \sigma_{-i}^*)) &= \{(3/8, \text{I}); (1/16, \text{II}); (1/8, \text{III}); (1/16, \text{IV}); (3/8, \text{V})\}, \\ L_i(\mu_i(\text{MF}, \text{MF}^b; \mathcal{M}_0, F, \sigma_{-i}^*)) &= \{(1/16, \text{I}); (3/8, \text{II}); (1/8, \text{III}); (3/8, \text{IV}); (1/16, \text{V})\}, \\ L_i(\mu_i(\text{MF}, \text{UF}; \mathcal{M}_0, F, \sigma_{-i}^*)) &= \{(7/32, \text{I}); (7/32, \text{II}); (1/8, \text{III}); (7/32, \text{IV}); (7/32, \text{V})\}, \\ L_i(\mu_i(\text{MF}, \text{SF}; \mathcal{M}_0, F, \sigma_{-i}^*)) &= \text{III}. \end{aligned}$$

The corresponding CPT values are as follows:

$$\begin{aligned} & V_i^{\text{MF}}(L_i(\mu_i(\text{MF}, \text{MF}^a; \mathcal{M}_0, F, \sigma_{-i}^*))) \\ &= 3.8w^+(20/32) + 2.016w^+(18/32) + 2.8w^+(14/32) + 5w^+(12/32) \\ &= 5.8243, \end{aligned}$$

$$\begin{aligned}
& V_i^{\text{MF}}(L_i(\mu_i(\text{MF}, \text{MF}^b; \mathcal{M}_0, F, \sigma_{-i}^*))) \\
&= 3.8w^+(30/32) + 2.016w^+(18/32) + 2.8w^+(14/32) + 5w^+(2/32) \\
&= 5.8243,
\end{aligned}$$

$$\begin{aligned}
& V_i^{\text{MF}}(L_i(\mu_i(\text{MF}, \text{UF}; \mathcal{M}_0, F, \sigma_{-i}^*))) \\
&= 3.8w^+(25/32) + 2.016w^+(18/32) + 2.8w^+(14/32) + 5w^+(7/32) \\
&= 5.6993,
\end{aligned}$$

and,

$$V_i^{\text{MF}}(L_i(\mu_i(\text{MF}, \text{SF}; \mathcal{M}_0, F, \sigma_{-i}^*))) = 5.816.$$

Thus $\sigma_i^*(\text{MF})$ has support on optimal signals, and hence is a best response. This completes the verification that σ^* is a Bayes-Nash equilibrium. With this, we end our example. \square

In the previous example, let us focus on the behavior of player i when she has type MF. For any mechanism with the signal sets for the players being $\Psi_1 = \Psi_2 = \{\text{MF}^a, \text{MF}^b, \text{UF}, \text{SF}\}$ as above (the mechanism $\mathcal{M}_0 = ((\Psi_i)_i, h_0)$ considered above is an instance of such a mechanism), the signals MF^a and MF^b allow this player to play so that the lotteries faced by her are $L'_i := L_i(\mu_i(\text{MF}, \text{MF}^a); \mathcal{M}, F, \sigma_{-i})$ and $L''_i := L_i(\mu_i(\text{MF}, \text{MF}^b); \mathcal{M}, F, \sigma_{-i})$ respectively, where F denotes the prior distribution on types (i.e. the product distribution with marginals given as in (6.3.3) above) and σ_{-i} denotes the strategy of the other player. The lotteries L'_i and L''_i are equally preferred by player i when she has type MF, and they are preferred over the lotteries corresponding to signaling UF or SF, when the mechanism is $\mathcal{M}_0 = ((\Psi_i)_i, h_0)$ as considered in Example 6.3.1, and the other player plays according to the strategy prescribed in (6.3.7). Under the equilibrium strategy σ_i^* , as defined in (6.3.7), when player i has type MF she signals MF^a or MF^b each with probability half.

We can think of player 1 as tossing a fair coin to decide whether to signal MF^a or MF^b when her type is MF, and similarly for player 2. The outcome of the coin toss is private knowledge to the player tossing the coin. The equilibrium strategies in (6.3.7) correspond to each player signaling UF or SF truthfully if that is her type, while if her type is MF then she signals MF^a or MF^b depending on the outcome of her coin toss. From our analysis in the above example, we have that these strategies form an F -Bayes-Nash equilibrium for this game and induce the allocation choice function f^* .

An alternate viewpoint is to think of the coins being tossed at the beginning as before, but now let us assume that the system operator observes the outcomes of both the coins. We continue to assume that each player does not observe the result of the coin toss of the other player. Suppose each player only has the option to signal from $\{\text{MF}, \text{UF}, \text{SF}\}$. The system operator collects these signals and implements a lottery on the allocation set according to the following rule: If player i signals UF or SF then let $\psi'_i = \text{UF}$ or $\psi'_i = \text{SF}$ respectively. If player i signals MF then, depending on the outcome of coin toss i , let $\psi'_i = \text{MF}^a$ or MF^b . The system operator implements $h_0(\psi'_1, \psi'_2)$. Now consider the strategy where each player

reports her type truthfully. We observe that this strategy is an F -Bayes-Nash equilibrium for this game and induces f^* .

Thus the issue with the revelation principle is superficial in the sense that the reason that it does not hold is not that player i does not wish to reveal her type, but rather that she would like to have an asymmetry in the information of the players. In the above example, this asymmetry comes from the coin tosses and, as seen in the latter viewpoint, these coin tosses can be thought of as shared between each individual player and the system operator, so one could even think of the coins as being tossed by system operator, with the result of each individual coin toss being shared with the respective player. To capture this intuition, we propose a framework where there is a mediator who sends messages to each individual player before collecting their signals. As we will see now, this way we can recover a form of the revelation principle.

6.4 Mediated Mechanisms and the Revelation Principle

We now lay out the framework for a mechanism with messages from the mediator, along the lines of the augmented framework for mechanism design motivated by the example above. Let Φ_i be a finite *message set* for each player i , with a typical element denoted by ϕ_i , and let $\Phi := \prod_i \Phi_i$. Let $D \in \Delta(\Phi)$ denote a *mediator distribution* from which the mediator draws a profile of messages $\phi := (\phi_1, \dots, \phi_n)$. Let $D_i \in \Delta(\Phi_i)$ denote the marginal of D on Φ_i . For any $\phi_i \in \text{supp } D_i$, let the conditional distribution be given by

$$D_{-i}(\phi_{-i}|\phi_i) := \frac{D(\phi_i, \phi_{-i})}{D_i(\phi_i)}, \text{ for all } \phi_{-i} \in \Phi_{-i}, \quad (6.4.1)$$

where $\phi_{-i} := (\phi_j)_{j \neq i}$ and $\Phi_{-i} := \prod_{j \neq i} \Phi_j$. Let Ψ_i be a finite set of signals as before. Let

$$h : \Phi \times \Psi \rightarrow \Delta(A) \quad (6.4.2)$$

be a *mediated allocation function*. The message sets $\Phi_i, i \in [n]$, a mediator distribution $D \in \Delta(\Phi)$, and a mediated allocation function h together constitute a *mediated mechanism*, denoted by

$$\mathcal{M} := ((\Phi_i)_{i \in [n]}, D, (\Psi_i)_{i \in [n]}, h). \quad (6.4.3)$$

The mediator first draws a profile of messages ϕ from the distribution D . Each player i observes her message ϕ_i , and then sends a signal ψ_i to the mediator. The mediator collects the signals from all the players and then chooses an allocation according to the probability distribution $h(\phi, \psi)$. A strategy for any player i is thus given by

$$\tau_i : \Phi_i \times \Theta_i \rightarrow \Delta(\Psi_i). \quad (6.4.4)$$

Let $\tau_i(\psi_i|\phi_i, \theta_i)$ denote the probability of signal ψ_i under the distribution $\tau_i(\phi_i, \theta_i)$. Let $\tau := (\tau_1, \dots, \tau_n)$ denote the profile of strategies. Suppose player i has received message

ϕ_i and has type θ_i (thus, $\phi_i \in \text{supp } D_i$, and $\theta_i \in \text{supp } F_i$), and she chooses to signal ψ_i (so $\psi_i \in \text{supp } \tau_i(\phi_i, \theta_i)$); then consider the probability distribution $\mu_i(\phi_i, \theta_i, \psi_i; \mathcal{M}, F, \tau_{-i}) \in \Delta(A)$ given by

$$\begin{aligned} \mu_i(\alpha|\phi_i, \theta_i, \psi_i; \mathcal{M}, F, \tau_{-i}) &:= \sum_{\phi_{-i}} D_{-i}(\phi_{-i}|\phi_i) \sum_{\theta_{-i}} F_{-i}(\theta_{-i}|\theta_i) \\ &\quad \times \sum_{\psi_{-i}} \prod_{j \neq i} \tau_j(\psi_j|\phi_j, \theta_j) h(\alpha|\phi, \psi). \end{aligned} \quad (6.4.5)$$

The *best response strategy set* $BR_i(\tau_{-i})$ of player i to a strategy profile τ_{-i} of her opponents consists of all strategies $\tau_i^* : \Phi_i \times \Theta_i \rightarrow \Delta(\Psi_i)$ such that

$$W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi_i; \mathcal{M}, F, \tau_{-i})) \geq W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi'_i; \mathcal{M}, F, \tau_{-i})), \quad (6.4.6)$$

for all $\phi_i \in \text{supp } D_i$, $\theta_i \in \text{supp } F_i$, $\psi_i \in \text{supp } \tau_i^*(\phi_i, \theta_i)$, $\psi'_i \in \Psi_i$.

A strategy profile τ^* is said to be an *F-Bayes-Nash equilibrium* for the environment \mathcal{E} with respect to the mediated mechanism \mathcal{M} if for each player i we have

$$\tau_i^* \in BR_i(\tau_{-i}^*). \quad (6.4.7)$$

We will say that an allocation choice function $f : \Theta \rightarrow \Delta(A)$ is implementable in *F-Bayes-Nash equilibrium* by a mediated mechanism if there exists a mediated mechanism \mathcal{M} and an *F-Bayes-Nash equilibrium* τ with respect to this mediated mechanism such that f is the induced allocation choice function, i.e. for all $\theta \in \text{supp } F$, $\alpha \in A$, we have

$$f(\alpha|\theta) = \sum_{\phi} D(\phi) \sum_{\psi} \left(\prod_i \tau_i(\psi_i|\phi_i, \theta_i) \right) h(\alpha|\phi, \psi). \quad (6.4.8)$$

A mediated mechanism $\mathcal{M} = ((\Phi_i)_{i \in [n]}, D, (\Psi_i)_{i \in [n]}, h)$ is called a *direct mediated mechanism* if $\Psi_i = \Theta_i$ for all i , and we write it as $\mathcal{M}^d = ((\Phi_i)_{i \in [n]}, D, (\Theta_i)_{i \in [n]}, h^d)$, where

$$h^d : \Phi \times \Theta \rightarrow \Delta(A)$$

is the corresponding *direct mediated allocation function*.

For a direct mediated mechanism, the truthful strategy τ_i^d for player i should satisfy $\tau_i^d(\phi_i, \theta_i) = \theta_i$, for all $\phi_i \in \Phi_i$, and $\theta_i \in \Theta_i$. Thus, if player i receives a message ϕ_i and has type θ_i , she reports her true type θ_i irrespective of her received message. In a way, the messages are present only to align the beliefs of the players appropriately so that truth-telling is an equilibrium strategy (depending on the type of equilibrium under consideration, i.e. Bayes-Nash, dominant, or belief-dominant equilibrium). Note that in the definition of the truthful strategy τ_i^d for player i we require $\tau_i^d(\phi_i, \theta_i) = \theta_i$, for all $\theta_i \in \Theta_i$ and $\phi_i \in \Phi_i$, and not just for $\theta_i \in \text{supp } F_i$ (when discussing an *F-Bayes-Nash equilibrium*) and $\phi_i \in \text{supp } D_i$. This is done to make the notion of a truthful strategy uniquely defined.

An allocation choice function f is said to be truthfully implementable in mediated F -Bayes-Nash equilibrium if there exists a direct mediated mechanism \mathcal{M}^d such that the truthful strategy profile τ^d is a mediated F -Bayes-Nash equilibrium and it implements f .

Let us now extend the notion of dominant equilibrium and belief-dominant equilibrium to the mediated mechanism framework. Let

$$\mu_i(\phi_i, \theta_i, \psi_i; \mathcal{M}, \psi_{-i}) := \sum_{\phi_{-i}} D_{-i}(\phi_{-i} | \phi_i) h(\phi, \psi), \quad (6.4.9)$$

denote the lottery faced by player i with type θ_i , who has received message ϕ_i (thus, $\phi_i \in \text{supp } D_i$) and believes that her opponents are going to report ψ_{-i} . Similarly, let

$$\mu_i(\phi_i, \theta_i, \psi_i; \mathcal{M}, G_{-i}) := \sum_{\phi_{-i}} D_{-i}(\phi_{-i} | \phi_i) \sum_{\psi_{-i}} G_{-i}(\psi_{-i}) h(\phi, \psi), \quad (6.4.10)$$

denote the lottery faced by player i with type θ_i , who has received message $\phi_i \in \text{supp } D_i$ and has belief $G_{-i} \in \Delta(\Psi_{-i})$ over her opponents' signal profiles. We define strategy τ_i^* to be dominant if, for all $\phi_i \in \text{supp } D_i$, $\theta_i \in \Theta_i$, $\psi_i \in \text{supp } \tau_i^*(\phi_i, \theta_i)$, $\psi'_i \in \Psi_i$, and $\psi_{-i} \in \Psi_{-i}$, we have

$$W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi_i; \mathcal{M}, \psi_{-i})) \geq W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi'_i; \mathcal{M}, \psi_{-i})). \quad (6.4.11)$$

Similarly, we define strategy τ_i^* to be belief-dominant if, for all $\phi_i \in \text{supp } D_i$, $\theta_i \in \Theta_i$, $\psi_i \in \text{supp } \tau_i^*(\phi_i, \theta_i)$, $\psi'_i \in \Psi_i$, and $G_{-i} \in \Delta(\Psi_{-i})$, we have

$$W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi_i; \mathcal{M}, G_{-i})) \geq W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi'_i; \mathcal{M}, G_{-i})). \quad (6.4.12)$$

An allocation choice function f is said to be *implementable in dominant equilibrium* by a mediated mechanism if there is a mediated mechanism \mathcal{M} and a dominant equilibrium τ (i.e. a strategy profile comprised of dominant strategies for the individual players) such that f is the allocation choice function induced by τ under \mathcal{M} , i.e. (6.4.8) holds for all $\theta \in \Theta$ and $\alpha \in A$. f is said to be *truthfully implementable in dominant equilibrium* by a direct mediated mechanism if there is a directed mediated mechanism \mathcal{M}^d such that the truthful strategy profile is a dominant equilibrium and induces f under \mathcal{M}^d . The notions of implementability by a mediated mechanism and truthful implementability by a direct mediated mechanism of an allocation choice function in belief-dominant equilibrium can be similarly defined.

If the message set Φ_i is a singleton for each player i , then we get back the original mechanism design framework. Thus, the mediated mechanism design framework defined above is a generalization of the mechanism design framework. This generalization allows us to establish a form of the revelation principle even when players have CPT preferences.

A special case of the mediated mechanism design framework is when the mediator message profile ϕ is publicly known. That is, each player knows the entire message profile instead of privately knowing only her own message. This would happen if $\Phi_i = \Phi_*$, for all $i \in [n]$, and D is a diagonal distribution, i.e. $D(\phi) = 0$ for all message profiles $\phi = (\phi_i)_{i \in [n]}$ such that

$\phi_i \neq \phi_j$ for some pair $i, j \in [n]$. Let Φ_* denote the common message set and let $D_* \in \Delta(\Phi_*)$ denote the mediator distribution on this set. Let

$$\mathcal{M}_* := (\Phi_*, D_*, (\Psi_i)_{i \in [n]}, h_*)$$

denote such a mediated mechanism with common messages, where now

$$h_* : \Phi_* \times \Psi \rightarrow \Delta(A).$$

We will call \mathcal{M}_* a *publicly mediated mechanism*. The notions of an allocation choice function being implementable in publicly mediated Bayes-Nash equilibrium, publicly mediated dominant equilibrium, or publicly mediated belief-dominant equilibrium can be defined similarly to the corresponding earlier definitions that were made for general message sets. The notions of an allocation choice function being truthfully implementable in direct publicly mediated Bayes-Nash equilibrium, direct publicly mediated dominant equilibrium, or direct publicly mediated belief-dominant equilibrium can also be defined similarly to the corresponding earlier definitions that were made for general message sets.

We are now in a position to state one of our main results.

Theorem 6.4.1 (Revelation Principle). *We have the following three versions of the revelation principle:*

- (i) *If an allocation choice function is implementable in Bayes-Nash equilibrium by a mediated mechanism, then it is also truthfully implementable in Bayes-Nash equilibrium by a direct mediated mechanism.*
- (ii) *If an allocation choice function is implementable in dominant equilibrium by a publicly mediated mechanism, then it is also truthfully implementable in dominant equilibrium by a direct publicly mediated mechanism.*
- (iii) *If an allocation choice function is implementable in belief-dominant equilibrium by a mediated (resp. publicly mediated) mechanism, then it is also truthfully implementable in belief-dominant equilibrium by a direct mediated (resp. direct publicly mediated) mechanism.*

We prove this theorem in Appendix 6.A. Theorem 6.4.1, in particular, implies that if an allocation choice function is implementable in Bayes-Nash equilibrium by a non-mediated mechanism then it is truthfully implementable in Bayes-Nash equilibrium by a direct mediated mechanism. Similarly, if an allocation choice function is implementable in dominant strategies (resp. belief-dominant strategies) by a non-mediated mechanism, then it is truthfully implementable in dominant strategies (resp. belief-dominant strategies) by a direct publicly mediated mechanism. Table 6.1 summarizes the different implementability settings under which the revelation principle does and does not hold. Example 6.3.1 shows that the revelation principle does not hold for the setting with Bayes-Nash equilibrium and non-mediated mechanism. In Example 6.4.2, we show that the revelation principle does not hold

	Non-mediated	Publicly Mediated	Mediated
Bayes-Nash Equilibrium	✗	✗	✓
Dominant Equilibrium	✗	✓	✗
Belief-dominant Equilibrium	✗	✓	✓

Table 6.1: Settings in which the revelation principle holds.

for the settings with dominant equilibrium or belief-dominant equilibrium and non-mediated mechanism. In Example 6.4.3, we show that the revelation principle does not hold for the settings with Bayes-Nash equilibrium and publicly mediated mechanism. Finally, in Example 6.4.4, we show that the revelation principle does not hold for the setting with dominant equilibrium and mediated mechanism.

Example 6.4.2. Consider the setting from Example 6.2.1 with two players. Recall that $\Theta_1 = \Theta_2 = \{\text{UP}, \text{DN}\}$, $A = \{a, b, c\}$, $\Gamma_1 = \{\text{I}, \text{II}, \text{III}, \text{IV}, \text{V}\}$, $\Gamma_2 = A$. The allocation-outcome mapping is given by the product distribution of the marginals ζ_1 and ζ_2 , given by $\zeta_1(a) = \{(1/2, \text{I}); (1/2, \text{V})\}$, $\zeta_1(b) = \{(1/2, \text{II}); (1/2, \text{IV})\}$, $\zeta_1(c) = \{(1, (\text{III}))\}$, and $\zeta_2(\alpha) = \alpha, \forall \alpha \in A$. The probability weighting functions for gains for the two players are

$$w_1^+(p) = \exp\{-(-\ln p)^{0.5}\}, w_2^+(p) = p,$$

for $p \in [0, 1]$ (see Figure 6.1). Let the value functions v_1 and v_2 be given by

v_1	I	II	III	IV	V
UP	$2x$	$x + 1$	1.99	1	0
DN	0	0	1	0	0

v_2	a	b	c
UP	1	0	2
DN	0	1	2

where $x := 1/w_1^+(0.5) = 2.2992$.

Consider a mechanism $\mathcal{M}_0 = \{(\Psi_1, \Psi_2), h_0\}$, where $\Psi_1 = \{a, b, c\}$, $\Psi_2 = \{\text{UP}, \text{DN}\}$, and

$$\begin{aligned} h_0(a, \psi_2) &= a, \\ h_0(b, \psi_2) &= b, \\ h_0(c, \psi_2) &= c, \end{aligned}$$

for all $\psi_2 \in \Psi_2$. The CPT values for player 1 having type UP for the lotteries over her outcomes corresponding to the different allocations are given by

$$\begin{aligned} V_1^{\text{UP}}(L_1(a)) &= 2xw_1^+(0.5) = 2, \\ V_1^{\text{UP}}(L_1(b)) &= w_1^+(1) + xw_1^+(0.5) = 2, \\ V_1^{\text{UP}}(L_1(c)) &= 1.99. \end{aligned}$$

Further, the CPT values for player 1 having type DN for the above lotteries are given by

$$V_1^{\text{DN}}(L_1(a)) = 0, \quad V_1^{\text{DN}}(L_1(b)) = 0, \quad V_1^{\text{DN}}(L_1(c)) = 1.$$

Since the allocation choice function h_0 does not depend on the signal of player 2, from the above the above calculations, we observe that the strategy σ_1 given by

$$\begin{aligned}\sigma_1(\cdot|0) &= \{(0.5, a); (0.5, b)\}, \\ \sigma_1(\cdot|1) &= c,\end{aligned}$$

is a dominant strategy. and a belief-dominant strategy. Let σ_2 be the truthful strategy for player 2. Again, since the allocation choice function h_0 does not depend on the signal of player 2, σ_2 is trivially a dominant strategy and a belief-dominant strategy. Thus $\sigma = (\sigma_1, \sigma_2)$ is a dominant equilibrium and a belief-dominant equilibrium. The corresponding social choice function f is given by

$$\begin{aligned}f(\text{UP}, \theta_2) &= \{(0.5, a); (0.5, b)\}, \\ f(\text{DN}, \theta_2) &= c.\end{aligned}$$

Thus, the allocation choice function f is implementable in dominant (resp. belief-dominant) equilibrium. Suppose there were a direct mechanism $\mathcal{M}_0^d = h_0^d$ that truthfully implements the allocation choice function f in dominant (resp. belief-dominant) equilibrium. Then, $h_0^d = f$. As observed in Example 6.2.1, the CPT value for player 1 having type UP for the lottery corresponding to $\{(0.5, a); (0.5, b)\}$ is

$$V_1^{\text{UP}}(L_1(\{(0.5, a); (0.5, b)\})) = 1.9851.$$

If player 1 has type UP and believes that player 2's type report is UP (or equivalently, any other distribution over player 2's type report), then player 1 would deviate from her truthful strategy and report DN instead, because it gives her a higher CPT value. Hence the truthful strategy σ_1^d is not a dominant (resp. belief-dominant) equilibrium for the direct mechanism \mathcal{M}_0^d . Thus f is not truthfully implementable in dominant (resp. belief-dominant) equilibrium by a direct mechanism. \square

We will now show that the revelation principle does not hold for the setting with Bayes-Nash equilibrium and publicly mediated mechanism. Let us first make an observation regarding the allocation choice functions that are truthfully implementable in F -Bayes-Nash equilibrium by a direct publicly mediated mechanism. Let f be an allocation choice function that is truthfully implementable in F -Bayes-Nash equilibrium by a direct publicly mediated mechanism

$$\mathcal{M}_*^d = (\Phi_*, D_*, (\Theta_i)_{i \in [n]}, h_*^d), \tag{6.4.13}$$

where

$$h_*^d : \Phi_* \times \Theta \rightarrow \Delta(A), \tag{6.4.14}$$

is the direct mediated allocation function for this direct publicly mediated mechanism. Since truthful strategies τ^d are an F -Bayes-Nash equilibrium, for each $\phi_* \in \text{supp } D_*$, we have

$$W_i^{\theta_i}(\mu_i(\phi_*, \theta_i, \theta_i; \mathcal{M}_*^d, F, \tau_{-i}^d)) \geq W_i^{\theta_i}(\mu_i(\phi_*, \theta_i, \tilde{\theta}_i; \mathcal{M}_*^d, F, \tau_{-i}^d)), \tag{6.4.15}$$

for all $\theta_i \in \text{supp } F_i, \tilde{\theta}_i \in \Theta_i, i \in [n]$, where

$$\mu_i(\phi_*, \theta_i, \tilde{\theta}_i; \mathcal{M}_*^d, F, \tau_{-i}^d) = \sum_{\theta_{-i}} F_{-i}(\theta_{-i} | \theta_i) h_*^d(\phi_*, \tilde{\theta}_i, \theta_{-i}),$$

is the lottery induced on the allocations for player i receiving message ϕ_* , having type θ_i , and deciding to report type $\tilde{\theta}_i$. Now, fix $\phi_* \in \Phi_*$ with $D_*(\phi_*) > 0$, and consider a non-mediated direct mechanism $\mathcal{M}_0^d := ((\Theta_i)_{i \in [n]}, h_0^d)$, with its direct allocation function being $h_0^d(\cdot) := h_*^d(\phi_*, \cdot) : \Theta \rightarrow \Delta(A)$. It follows from (6.4.15) that truthful strategies corresponding to mechanism \mathcal{M}_0^d form an F -Bayes-Nash equilibrium. Thus, we note that $h_*^d(\phi_*, \cdot)$ is the allocation function truthfully implemented by the non-mediated direct mechanism \mathcal{M}_0^d . Since mechanism \mathcal{M}_*^d truthfully implements the allocation function f in F -Bayes-Nash equilibrium, we have that

$$f(\theta) = \sum_{\phi_*} D_*(\phi_*) h_*^d(\phi_*, \theta),$$

for all $\theta \in \text{supp } F$. From these two observations, we conclude that if f is an allocation choice function that is truthfully implementable in F -Bayes-Nash equilibrium by a direct publicly mediated mechanism, then f is a convex combination of allocation choice functions each of which is truthfully implementable in F -Bayes-Nash equilibrium by a non-mediated direct mechanism. It is easy to see that the converse of this statement is also true.

In the following example, we will use this observation to establish that the revelation principle does not hold for the setting with Bayes-Nash equilibrium and publicly mediated mechanism.

Example 6.4.3. Let there be two players, i.e. $n = 2$. Let $\Theta_1 = \Theta_2 = \{\text{UP}, \text{DN}\}$. Let $\Gamma_1 = \Gamma_2 = \{\text{I}, \text{II}, \text{III}, \text{IV}, \text{V}\}$. Let the value function v_1 for player 1 be as shown below

v_1	I	II	III	IV	V
UP	80	57	34	17	0
DN	0	0	100	0	0

and let the value function v_2 for player 2 be as shown below

v_2	I	II	III	IV	V
UP	-79	-56	-33	-17	0
DN	0	0	100	0	0

Let the probability weighting functions for both the players, for both types, for gains and losses, be given by the following piecewise linear function:

$$w_1^\pm(p) = w_2^\pm(p) = w(p) = \begin{cases} (8/7)p, & \text{for } 0 \leq p < (7/32), \\ (1/4) + (2/3)(p - 7/32), & \text{for } (7/32) \leq p < 25/32, \\ (5/8) + (12/7)(p - 25/32), & \text{for } (25/32) \leq p < 1, \end{cases}$$

(See the probability weighting function for gains in Figure 6.1.) Let the prior distribution F be such that the types of the players are independently sampled with probabilities,

$$\mathbb{P}(\text{UP}) = 3/4, \mathbb{P}(\text{DN}) = 1/4. \quad (6.4.16)$$

Let $A = \{a, b, c\}$. Let

$$\begin{aligned} \zeta(a) &= \{(1/2, (\text{I}, \text{I})); (1/2, (\text{V}, \text{V}))\}, \\ \zeta(b) &= \{(1/2, (\text{II}, \text{II})); (1/2, (\text{IV}, \text{IV}))\}, \\ \zeta(c) &= (\text{III}, \text{III}). \end{aligned}$$

Consider the allocation choice function f^* given by

$$\begin{aligned} f^*(\text{DN}, \theta_2) &= f^*(\theta_1, \text{DN}) = c, \quad \forall \theta_1 \in \Theta_1, \theta_2 \in \Theta_2 \\ f^*(\text{UP}, \text{UP}) &= \{(1/2, a); (1/2, b)\}. \end{aligned}$$

We will now show that f^* is implementable in F -Bayes-Nash equilibrium by a publicly mediated mechanism. In fact, we will show that f^* is implementable in F -Bayes-Nash equilibrium by a non-mediated mechanism. We will then show that f^* cannot be a convex combination of allocation choice functions each of which is truthfully implementable by a non-mediated direct mechanism. This will give us that f^* is not truthfully implementable in F -Bayes-Nash equilibrium by a direct publicly mediated mechanism. We will then conclude that the revelation principle does not hold for the setting with Bayes-Nash equilibrium and publicly mediated mechanism.

Consider the mechanism $\mathcal{M}_0 = ((\Psi_i)_{i \in [n]}, h_0)$, where $\Psi_1 = \{\text{UP}^a, \text{UP}^b, \text{DN}\}$, $\Psi_2 = \{\text{UP}, \text{DN}\}$, and the allocation function h_0 is given by

$$\begin{aligned} h_0(\text{DN}, \psi_2) &= h_0(\psi_1, \text{DN}) = c, \quad \forall \psi_1 \in \Psi_1, \psi_2 \in \Psi_2, \\ h_0(\text{UP}^a, \text{UP}) &= a, \\ h_0(\text{UP}^b, \text{UP}) &= b. \end{aligned}$$

Consider the strategy σ_1 for player 1 given by

$$\begin{aligned} \sigma_1(\text{UP}) &= \{(1/2, \text{UP}^a); (1/2, \text{UP}^b)\}, \\ \sigma_1(\text{DN}) &= \text{DN}, \end{aligned}$$

and the strategy σ_2 for player 2 given by

$$\begin{aligned} \sigma_2(\text{UP}) &= \text{UP}, \\ \sigma_2(\text{DN}) &= \text{DN}. \end{aligned}$$

It is easy to see that this induces the allocation choice function f^* .

We will now verify that σ is an F -Bayes-Nash equilibrium for \mathcal{M}_0 . If player 1 has type UP, then the CPT values of the lotteries faced by her corresponding to her signals are as follows:

$$\begin{aligned} W_1^{\text{UP}}(\mu_1(\text{UP}, \text{UP}^a; \mathcal{M}_0, F, \sigma_{-1})) &= V_1^{\text{UP}}(\{(3/8, \text{I}); (0, \text{II}); (1/4, \text{III}); (0, \text{IV}); (3/8, \text{V})\}) \\ &= 46w(3/8) + 34w(5/8) \\ &= 34. \end{aligned}$$

$$\begin{aligned} W_1^{\text{UP}}(\mu_1(\text{UP}, \text{UP}^b; \mathcal{M}_0, F, \sigma_{-1})) &= V_1^{\text{UP}}(\{(0, \text{I}); (3/8, \text{II}); (1/4, \text{III}); (3/8, \text{IV}); (0, \text{V})\}) \\ &= 23w(3/8) + 17w(5/8) + 17 \\ &= 34. \end{aligned}$$

$$W_1^{\text{UP}}(\mu_1(\text{UP}, \text{DN}; \mathcal{M}_0, F, \sigma_{-1})) = V_1^{\text{UP}}(\text{III}) = 34.$$

Thus player 1 is indifferent between all signals when she has type UP and so the strategy of signaling $\sigma_1(\text{UP}) = \{(1/2, \text{UP}^a); (1/2, \text{UP}^b)\}$ is optimal for her.

If player 1 has type DN, then III is the best outcome and she receives this lottery if she signals DN. Thus DN dominates any other strategy, in particular, signaling UP^a or UP^b .

If player 2 has type UP, then the CPT values of the lotteries faced by her corresponding to her signals are as follows:

$$\begin{aligned} W_2^{\text{UP}}(\mu_1(\text{UP}, \text{UP}; \mathcal{M}_0, F, \sigma_{-2})) &= V_1^{\text{UP}}(\{(3/16, \text{I}); (3/16, \text{II}); (1/4, \text{III}); (3/16, \text{IV}); (3/16, \text{V})\}) \\ &= -23w(3/16) - 23w(3/8) - 16w(5/8) - 17w(13/16) \\ &= -32.94. \end{aligned}$$

$$W_2^{\text{UP}}(\mu_1(\text{UP}, \text{DN}; \mathcal{M}_0, F, \sigma_{-2})) = V_1^{\text{UP}}(\text{III}) = -33.$$

Hence the strategy of signaling $\sigma_2(\text{UP}) = \text{UP}$ is optimal for player 2 when she has type UP.

If player 2 has type DN, then III is the best outcome and she receives this lottery if she signals DN. Thus DN dominates any other strategy, in particular, signaling UP.

This shows that σ is an F -Bayes-Nash equilibrium for \mathcal{M}_0 , and hence establishes that f^* is implementable in F -Bayes-Nash equilibrium by a non-mediated mechanism.

Suppose f^* were a convex combination of allocation choice functions each of which is truthfully implementable by a non-mediated direct mechanism. Let f be one of the allocation choice functions in this convex combination. Since $f^*(\text{DN}, \theta_2) = f^*(\theta_1, \text{DN}) = c$ for all θ_1, θ_2 , and since $\{c\}$ is an extreme point of the simplex $\Delta(A)$, we get that

$$f(\text{DN}, \theta_2) = f(\theta_1, \text{DN}) = c, \quad \forall \theta_1 \in \Theta_1, \theta_2 \in \Theta_2. \quad (6.4.17)$$

Similarly, since $f^*(\text{UP}, \text{UP})$ lies on the line joining the vertices $\{a\}$ and $\{b\}$ of the simplex $\Delta(A)$, we get that

$$f(\text{UP}, \text{UP}) = \{(x, a); (1-x, b)\}, \quad (6.4.18)$$




Figure 6.1: Plot of expression $E_3(x)$ in Example 6.4.3.

where $0 \leq x \leq 1$.

Let f be truthfully implementable in F -Bayes-Nash equilibrium by the non-mediated direct mechanism $\mathcal{M}_0^d = h_0^d$. Then $h_0^d = f$. If player 1 has type UP, then the lottery faced by her if she reports UP is given by

$$L_1(\mu_1(\text{UP}, \text{UP}; \mathcal{M}_0^d, F, \sigma_{-1}^d)) = \{(3x/8, \text{I}); (3(1-x)/8, \text{II}); (1/4, \text{III}); (3(1-x)/8, \text{IV}); (3x/8, \text{V})\},$$

where $\sigma_{-1}^d = \sigma_2^d$ is the truthful strategy of player 2. Let

$$E_3(x) := 23w \left(\frac{3x}{8} \right) + 23w \left(\frac{3}{8} \right) + 17w \left(\frac{5}{8} \right) + 17w \left(\frac{1-3x}{8} \right),$$

for $x \in [0, 1]$. We observe that $E_3(x)$ is maximum at $x = 0$ and $x = 1$, and for all $x \in (0, 1)$, $E_3(x) < 34$. (See Figure 6.1.)

Now, unless $x = 0$ or $x = 1$, player 1 will defect from the truthful strategy and report DN when her true type is UP, because if she does so the allocation c will be implemented by the system operator, which results in her outcome III, hence giving her a value of 34. Thus, $x = 0$ or $x = 1$.

If player 2 has type UP, then the lottery faced by her if she reports UP is given by

$$L_2(\mu_2(\text{UP}, \text{UP}; \mathcal{M}_0^d, F, \sigma_{-2}^d)) = \{(3x/8, \text{I}); (3(1-x)/8, \text{II}); (1/4, \text{III}); (3(1-x)/8, \text{IV}); (3x/8, \text{V})\},$$

where $\sigma_{-2}^d = \sigma_1^d$ is the truthful strategy of player 1. If $x = 0$, then the CPT value for player 2 is given by

$$\begin{aligned} V_2^{\text{UP}}(\{(0, \text{I}); (3/8, \text{II}); (1/4, \text{III}); (3/8, \text{IV}); (0, \text{V})\}) \\ &= -23w(3/8) - 16w(5/8) - 17 \\ &= -33.48. \end{aligned}$$

If $x = 1$, then the CPT value for player 2 is given by

$$\begin{aligned} & V_2^{\text{UP}}(\{(3/8, \text{I}); (0, \text{II}); (1/4, \text{III}); (0, \text{IV}); (3/8, \text{V})\}) \\ &= -46w(3/8) - 33w(5/8) \\ &= -33.48. \end{aligned}$$

Now, if $x = 0$ or $x = 1$, player 2 will defect from the truthful strategy and report DN when her true type is UP, because if she does so the allocation c will be implemented by the system operator, which results in her outcome III, hence giving her a value of -33 . Thus x cannot be 0 or 1, leading to a contradiction. Thus, f^* cannot be a convex combination of allocation choice functions each of which is truthfully implementable by a non-mediated direct mechanism. This completes the argument. \square

Example 6.4.4. We will now show that in the setting with dominant equilibrium and mediated mechanism, the revelation principle does not hold. This happens even when the players have EUT preferences. Let there be two players, i.e. $n = 2$. Let $\Theta_1 = \{\text{UP}, \text{DN}\}$, $\Theta_2 = \{\text{UN}\}$. Let $\Gamma_1 = \Gamma_2 = A = \{\text{I}, \text{II}, \text{III}\}$. Let the value functions v_1 and v_2 be given by

v_1	I	II	III
UP	5	-10	10
DN	0	0	10

v_2	I	II	III
UN	5	-10	-5

Both the players have EUT preferences (i.e, their probability weighting functions are identity functions).

Consider the allocation choice function f given by

$$f(\text{UP}, \text{UN}) = \text{I} \text{ and } f(\text{DN}, \text{UN}) = \text{III}.$$

We claim that this allocation choice function cannot be implemented by any direct mediated mechanism in truthful dominant equilibrium. To see this, suppose there exists a direct mediated mechanism $\mathcal{M}^d = ((\Phi_1, \Phi_2), D, h)$ that truthfully implements the allocation choice function f in dominant equilibrium. Consider the lottery

$$\mu_1(\phi_1, \theta_1, \theta'_1; \mathcal{M}^d, \text{UN}) = \sum_{\phi_2} D_{-1}(\phi_2 | \phi_1) h(\phi, \theta'_1, \text{UN}),$$

for player 1 when she receives message ϕ_1 , has type θ_1 , chooses to report θ'_1 (and trivially believes that player 2 reports type UN since that is the only type for player 2). Note that $\mu_1(\phi_1, \theta_1, \theta'_1; \mathcal{M}^d, \text{UN})$ does not depend on the type θ_1 of player 1. For any message ϕ_1 , consider the function $\pi_1(\phi_1, \cdot) : \Theta \rightarrow \Delta(A)$ given by

$$\pi_1(\phi_1, \theta_1, \text{UN}) := \mu_1(\phi_1, \theta_1, \theta_1; \mathcal{M}^d, \text{UN}).$$

Note that $\pi_1(\phi_1, \cdot)$ is an allocation choice function. Since \mathcal{M}^d truthfully implements f , we have

$$f(\theta_1, \text{UN}) = \sum_{\phi_1} D_1(\phi_1) \pi_1(\phi_1, \theta_1, \text{UN}),$$

for $\theta_1 = \text{UP}$ and DN . Thus, the allocation choice function f is a convex combination of the functions $\{\pi_1(\phi_1, \cdot)\}_{\phi_1}$. Since f is an extreme point of the convex set of all allocation choice functions, we get that $\pi_1(\phi_1, \cdot) = f(\cdot)$ for all ϕ_1 . Since truthful strategies are assumed to form a dominant equilibrium, we should have

$$W_1^{\theta_1}(\mu_1(\phi_1, \theta_1, \theta_1; \mathcal{M}, \text{UN})) \geq W_1^{\theta_1}(\mu_1(\phi_1, \theta_1, \theta'_1; \mathcal{M}, \text{UN})),$$

for $\theta_1 = \text{UP}$ and DN . However,

$$\begin{aligned} W_1^{\text{UP}}(\mu_1(\phi_1, \text{UP}, \text{UP}; \mathcal{M}, \text{UN})) &= W_1^{\text{UP}}(\pi_1(\phi_1, \text{UP}, \text{UN})) \\ &= W_1^{\text{UP}}(f(\text{UP}, \text{UN})) \\ &= W_1^{\text{UP}}(\text{I}) = 5 < 10 = W_1^{\text{UP}}(\text{III}) \\ &= W_1^{\text{UP}}(f(\text{DN}, \text{UN})) \\ &= W_1^{\text{UP}}(\pi_1(\phi_1, \text{DN}, \text{UN})) \\ &= W_1^{\text{UP}}(\mu_1(\phi_1, \text{UP}, \text{DN}; \mathcal{M}, \text{UN})). \end{aligned}$$

This is a contradiction, and hence, we conclude that the allocation choice function f is not truthfully implementable by a direct mediated mechanism in dominant equilibrium.

We will now show that the allocation choice function f can be implemented by mediated mechanism in dominant equilibrium if we use the message sets $\Phi_1 = \{C\}$, $\Phi_2 = \{L, R\}$ and the signal sets $\Psi_1 = \{\text{UP}, \text{DN}\}$, $\Psi_2 = \{L, R\}$. Let the mediator distribution be given by $D(C, L) = D(C, R) = 1/2$. Let the mediated allocation function $h : \Phi \times \Psi \rightarrow \Delta(A)$ be given by

$$\begin{aligned} h(C, L, \text{UP}, L) &= \text{I}, \\ h(C, L, \text{UP}, R) &= \text{I}, \\ h(C, L, \text{DN}, L) &= \text{III}, \\ h(C, L, \text{DN}, R) &= \text{II}, \\ h(C, R, \text{UP}, L) &= \text{I}, \\ h(C, R, \text{UP}, R) &= \text{I}, \\ h(C, R, \text{DN}, L) &= \text{II}, \\ h(C, R, \text{DN}, R) &= \text{III}. \end{aligned}$$

Let $\mathcal{M} = ((C, (L, R)), D, h)$ be the mediated mechanism. Now, consider the strategies $\sigma_1 : \Phi_1 \times \Theta_1 \rightarrow \Delta(\Psi_1)$ and $\sigma_2 : \Phi_2 \times \Theta_2 \rightarrow \Delta(\Psi_2)$ given by

$$\begin{aligned} \sigma_1(C, \text{UP}) &= \text{UP} \text{ and } \sigma_1(C, \text{DN}) = \text{DN}, \\ \sigma_2(L, \text{UN}) &= L \text{ and } \sigma_2(R, \text{UN}) = R. \end{aligned}$$

We will now verify that these strategies form a dominant equilibrium and implement the allocation choice function f . Let us first verify that σ_1 is a dominant strategy. We have,

$$\begin{aligned} W_1^{\text{UP}}(\mu_1(C, \text{UP}, \text{UP}; \mathcal{M}, L)) &= 0.5v_1^{\text{UP}}(h(C, L, \text{UP}, L)) + 0.5v_1^{\text{UP}}(h(C, R, \text{UP}, L)) = 5, \\ W_1^{\text{UP}}(\mu_1(C, \text{UP}, \text{DN}; \mathcal{M}, L)) &= 0.5v_1^{\text{UP}}(h(C, L, \text{DN}, L)) + 0.5v_1^{\text{UP}}(h(C, R, \text{DN}, L)) = 0, \\ W_1^{\text{UP}}(\mu_1(C, \text{UP}, \text{UP}; \mathcal{M}, R)) &= 0.5v_1^{\text{UP}}(h(C, L, \text{UP}, R)) + 0.5v_1^{\text{UP}}(h(C, R, \text{UP}, R)) = 5, \\ W_1^{\text{UP}}(\mu_1(C, \text{UP}, \text{DN}; \mathcal{M}, R)) &= 0.5v_1^{\text{UP}}(h(C, L, \text{DN}, R)) + 0.5v_1^{\text{UP}}(h(C, R, \text{DN}, R)) = 0. \end{aligned}$$

Thus, when player 1 has type UP, it is in her best interest to report UP in both the cases corresponding to her belief about the report by player 2, namely, L and R . On the other hand, when player 1 has type DN, we have

$$\begin{aligned} W_1^{\text{DN}}(\mu_1(C, \text{DN}, \text{UP}; \mathcal{M}, L)) &= 0.5v_1^{\text{DN}}(h(C, L, \text{UP}, L)) + 0.5v_1^{\text{DN}}(h(C, R, \text{UP}, L)) = 0, \\ W_1^{\text{DN}}(\mu_1(C, \text{DN}, \text{DN}; \mathcal{M}, L)) &= 0.5v_1^{\text{DN}}(h(C, L, \text{DN}, L)) + 0.5v_1^{\text{DN}}(h(C, R, \text{DN}, L)) = 5, \\ W_1^{\text{DN}}(\mu_1(C, \text{DN}, \text{UP}; \mathcal{M}, R)) &= 0.5v_1^{\text{DN}}(h(C, L, \text{UP}, R)) + 0.5v_1^{\text{DN}}(h(C, R, \text{UP}, R)) = 0, \\ W_1^{\text{DN}}(\mu_1(C, \text{DN}, \text{DN}; \mathcal{M}, R)) &= 0.5v_1^{\text{DN}}(h(C, L, \text{DN}, R)) + 0.5v_1^{\text{DN}}(h(C, R, \text{DN}, R)) = 5. \end{aligned}$$

Thus, when player 1 has type DN, it is in her best interest to report DN in both the cases corresponding to her belief about the report by player 2, namely, L and R . Hence, σ_1 is a dominant strategy for player 1. For player 2, we have

$$\begin{aligned} W_2^{\text{UN}}(\mu_2(L, \text{UN}, L; \mathcal{M}, \text{UP})) &= v_2^{\text{UN}}(h(C, L, \text{UP}, L)) = 5, \\ W_2^{\text{UN}}(\mu_2(L, \text{UN}, R; \mathcal{M}, \text{UP})) &= v_2^{\text{UN}}(h(C, L, \text{UP}, R)) = 5, \\ W_2^{\text{UN}}(\mu_2(L, \text{UN}, L; \mathcal{M}, \text{DN})) &= v_2^{\text{UN}}(h(C, L, \text{DN}, L)) = -5, \\ W_2^{\text{UN}}(\mu_2(L, \text{UN}, R; \mathcal{M}, \text{DN})) &= v_2^{\text{UN}}(h(C, L, \text{DN}, R)) = -10. \end{aligned}$$

Hence, σ_2 is a dominant strategy for player 1.

Thus, σ is a dominant equilibrium. It is easy to verify that it implements the allocation choice function f for the mechanism \mathcal{M} . We thus conclude that the revelation principle does not hold in the setting with dominant equilibrium and mediated mechanism. \square

6.5 Summary

In this chapter, we considered mechanism design for CPT players and in this process we discovered several important concepts that have gone unnoticed in the classical setting with EUT players. Namely, we saw that it is important to treat the allocation set and the outcome set of each player separately, both from theoretical as well as behavioral point of view. We also saw that we need to be careful while considering the notion of dominant strategies with CPT players and we need to treat the two notions of dominant equilibrium and belief-dominant equilibrium separately. Next, we saw that the generalized framework of mediated mechanisms recovers the coveted revelation principle.

Notice that in Table 6.1, we have a tick mark in each row. This tells us that anything that can be implemented by a non-mediated mechanism (possibly with general signals) can be truthfully implemented by a direct (publicly or general) mediated mechanism. Indeed, if a social choice function is implementable in Bayes-Nash equilibrium by a non-mediated equilibrium, then it is truthfully implementable in Bayes-Nash equilibrium by a direct mediated mechanism. Further, the setting for implementability in Bayes-Nash equilibrium by mediated mechanisms is favorable in the sense that the revelation principle holds here and we can restrict our attention to truthful implementability by direct mediated mechanisms. In the setting for implementability in dominant equilibrium, if a social choice function is implementable in dominant equilibrium by a non-mediated equilibrium, then it is truthfully implementable in dominant equilibrium by a direct publicly mediated mechanism. And, the setting for implementability in dominant equilibrium by publicly mediated mechanisms is favorable in the sense that the revelation principle holds here and we can restrict our attention to truthful implementability by direct publicly mediated mechanisms. Finally, in the setting of implementability in belief-dominant equilibrium, if a social choice function is implementable in belief-dominant equilibrium by a non-mediated equilibrium, then it is truthfully implementable in belief-dominant equilibrium by a direct publicly mediated mechanism. Besides, the revelation principle holds for both publicly mediated and mediated mechanisms in the setting of belief-dominant equilibrium.

It is worthwhile to repeat the importance of truthful implementability of social choice functions by direct (mediated) mechanisms, namely, we can restrict our attention to truthful strategies. We will now see some of the benefits of the revelation principle and truthful strategies that make applications of mechanism design practical for large-scale implementation with participants who can be both, humans and machines.

Generally in the settings where agents exhibit deviations from expected utility behavior, one would expect that the participating agents do not possess large computational power. Hence, truthful strategies are especially suitable for such settings in contrast to the more complicated strategies that are permitted by the concept of Bayes-Nash equilibrium. On the other hand, if our participating agents do not possess large computational power, then it is natural to question if they have the ability to exhibit strategic behavior, in particular the requirement that the strategies form a Bayes-Nash equilibrium (or dominant equilibrium or belief-dominant equilibrium). However, there can also be agents in the system who do possess large computational power. Indeed, most of the systems such as online auctions and marketplaces or networked-systems such as transportation networks, Internet routing networks, etc. are comprised of players having varying degrees of computational and strategic abilities. For example, a firm participating in an online marketplace has the resources to estimate the common prior and other players' strategies through extensive data collection, and thus can develop optimal strategies. On the other hand, individual agents participating in the same system often lack such resources. When truthful strategies are in equilibrium, we get the best of both the worlds – it is easy for the players with limited resources to implement optimal strategies and at the same time there is no incentive for the players with large resources to deviate from these strategies.

Consider the setting when players have independent types, i.e. the common prior F on the type profiles has a product distribution $F = \prod_i F_i$. Let $\mathcal{M}^d = ((\Phi_i)_{i \in [n]}, D, (\Theta_i)_{i \in [n]}, h^d)$ be a direct mediated mechanism in such a setting. We note that the lottery induced on the outcome set of player i when she receives a message ϕ_i , has type θ_i , and decides to report $\tilde{\theta}$, is independent of her own type θ_i . This is because her belief $F_{-i}(\cdot | \theta_i)$ on the type profiles of her opponents is independent of her type θ_i . With an abuse of notation, let us denote this belief by $F_{-i} \in \Delta(\Theta_{-i})$. Then the lottery induced on the outcome set of player i when she receives a message $\phi_i \in \text{supp } D_i$, and decides to report $\tilde{\theta}$, is given by

$$L_i^{\phi_i, \tilde{\theta}_i}(\gamma_i) := \sum_{\phi_{-i}} D_{-i}(\phi_{-i} | \phi_i) \sum_{\theta_{-i}} F_{-i}(\theta_{-i}) \sum_{\alpha} h^d(\alpha | \phi, \tilde{\theta}, \theta_{-i}) \zeta_i(\gamma_i | \alpha), \quad \gamma_i \in \Gamma_i.$$

We will now interpret the message profile as determining the menu of options to be presented to each player. For example, if the message profile $\phi \in \Phi$ is drawn from the distribution D , then player i would be presented with the menu comprised of lotteries, one for each type $\tilde{\theta}_i \in \Theta_i$ of the player. Let

$$\mathcal{L}_i(\phi_i) := \{L_i^{\phi_i, \tilde{\theta}_i}\}_{\tilde{\theta}_i \in \Theta_i},$$

denote the list of lotteries presented to player i when her message is $\phi_i \in \Phi_i$. Depending on the player's type, she chooses the lottery that gives her maximum CPT value. If truthful strategies form an F -Bayes-Nash equilibrium, then the lottery $L_i^{\phi_i, \theta_i}$ is indeed the best option for a player with type θ_i .

In several practical situations, the players are unaware of the type sets of other players $\Theta_j, j \neq i$, the allocation set A , the allocation-outcome mapping ζ , and the common prior F . It might also be preferable to relieve the players from the burden of knowing the message sets and the mediator distribution D . Note that the system operator has enough knowledge to construct the list of lotteries $\mathcal{L}_i(\phi_i)$ for each player i based on her sampled message ϕ_i . Now, using the knowledge of her own type θ_i , namely her preferences on the lotteries over her outcome set, player i can select the lottery that is optimal for her from the list $\mathcal{L}_i(\phi_i)$. This provides a way to operate the mechanism \mathcal{M}^d under reasonable assumptions on the players' information.

Further, it is beneficial to limit the complexity of the list $\mathcal{L}_i(\phi_i)$ presented to the players. A way to do this would be to limit the size of the list and the complexity of each individual lottery in the list. The complexity of each individual lottery can be restricted, for example, by limiting the size of the outcome set Γ_i and by restricting the probabilities of each outcome to belong to a grid $\{k/K : 0 \leq k \leq K\}$, where $K > 0$ determines the granularity of the grid. Our framework with separate allocation and outcome sets is helpful in imposing restrictions on the size of the outcome set Γ_i . Subsequently, for any lottery $L_i \in \Delta(\Gamma_i)$, we can find an approximate lottery $\tilde{L}_i = \{(p_i(\gamma_i), \gamma_i)\}_{\gamma_i \in \Gamma_i}$ such that $p_i(\gamma_i) \in \{k/K : 0 \leq k \leq K\}$ for all $\gamma_i \in \Gamma_i$.

On the other hand, the size of the list $\mathcal{L}_i(\phi_i)$ is same as the size of the type set Θ_i in the worst case. This could make things practically infeasible. For example, when considering type spaces comprised of general CPT preferences, it might be impossible in practice to

elicit the probability weighting functions from the agents. Restricting the type space can lead to inefficient social choice functions. The mediated mechanism design framework could allow us to limit the size of menu options and at the same time have diversity in the social choice function across different types of the players, facilitated by the messaging stage. Such multiple communication rounds have been studied under EUT and there is an extensive literature concerning the communication requirements in mechanism design. (See [86] and the references therein. See also the literature on computational mechanism design [36].) Given that the non-EUT preferences can reliably be applied only to non-dynamic decision-making, we are especially interested in mechanisms that have a single stage of mediator messages to which the participating agents respond optimally by choosing their best option. It would be interesting to study the design of mechanisms that optimally elicit CPT preferences under communication restrictions such as limiting the size of the menu options. For example, we could consider mechanism designs where the mediated allocation function h^d for a direct mediated mechanism has to satisfy $|\{\mathcal{L}_i(\phi_i)\}| \leq B$, for all messages ϕ_i , for some bound B .

In this chapter, we focused on the mechanism design framework and the revelation principle for agents having CPT preferences. It is just the first step towards mechanism design for non-EUT players, with several interesting directions for future work. In the next chapter, we will discuss some of these directions.

Appendix

6.A Proof of the Revelation Principle

We will first consider the revelation principle in the setting of mediated mechanisms. This corresponds to statement (i) and a part of statement (ii) of Theorem 6.4.1. In this setting we will show that if an allocation choice function f is implementable in Bayes-Nash equilibrium (resp. belief-dominant equilibrium) by a mediated mechanism then it is truthfully implementable in Bayes-Nash equilibrium (resp. belief-dominant equilibrium) by a direct mediated mechanism. We will then consider the setting of publicly mediated mechanisms and show that if an allocation choice function f is implementable in dominant equilibrium (resp. belief-dominant equilibrium) by a publicly mediated mechanism then it is truthfully implementable in dominant equilibrium (resp. belief-dominant equilibrium) by a direct publicly mediated mechanism. This will complete the proof of statement (ii) and the remaining part of statement (iii) of Theorem 6.4.1.

For the first setting, let

$$\mathcal{M} = ((\Phi_i)_{i \in [n]}, D, (\Psi_i)_{i \in [n]}, h),$$

be a mediated mechanism and let τ be a strategy profile that induces f for this mechanism. Consider now the direct mediated mechanism

$$\mathcal{M}^d = ((\Phi'_i)_{i \in [n]}, D', (\Theta_i)_{i \in [n]}, h^d),$$

where the message set is given by

$$\Phi'_i := \Phi_i \times (\Psi_i)^{\Theta_i}, \quad (6.A.1)$$

with a typical element denoted by

$$\phi'_i := (\phi_i, (\psi_i^{\theta'_i})_{\theta'_i \in \Theta_i}), \quad (6.A.2)$$

and the mediator distribution D' is given by

$$D'(\phi') := D(\phi) \prod_{i \in [n]} \prod_{\theta'_i \in \Theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \text{ for all } \phi' \in \Phi'. \quad (6.A.3)$$

The modified mediator messages and the mediator distribution can be interpreted as encapsulating the randomness in the strategies of the players for each of their types into their private messages.

We now observe that

$$D'_i(\phi'_i) = D_i(\phi_i) \prod_{\theta'_i \in \Theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i), \quad (6.A.4)$$

and

$$\sum_{\phi'_i \in \Phi'_i} D'_i(\phi'_i) = \sum_{\phi' \in \Phi'} D'(\phi') = 1.$$

Thus, $D' \in \Delta(\Phi')$ is indeed a valid distribution. Equation (6.A.4) can be formally proved as follows:

$$\begin{aligned} D'_i(\phi'_i) &= \sum_{\phi'_{-i} \in \Phi'_{-i}} D'(\phi'_i, \phi'_{-i}) \\ &= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\substack{(\psi_j^{\theta'_j})_{\theta'_j \in \Theta_j, j \neq i} \\ \in \prod_{j \neq i} (\Psi_j)^{\Theta_j}}} \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j} \tau_j(\psi_j^{\theta'_j} | \phi_j, \theta'_j) \right) \prod_{\theta'_i \in \Theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \\ &= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \prod_{\theta'_i \in \Theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j} \sum_{\psi_j^{\theta'_j} \in \Psi_j} \tau_j(\psi_j^{\theta'_j} | \phi_j, \theta'_j) \right) \\ &= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \prod_{\theta'_i \in \Theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j} 1 \right) \end{aligned}$$

$$= D_i(\phi_i) \prod_{\theta'_i \in \Theta_i} \tau_i \left(\psi_i^{\theta'_i} | \phi_i, \theta'_i \right).$$

Let the direct mediated allocation function be given by

$$h^d(\phi', \theta') := h \left(\phi, \left(\psi_i^{\theta'_i} \right)_{i \in [n]} \right) \text{ for all } \phi' \in \Phi', \theta' \in \Theta. \quad (6.A.5)$$

Note that the construction of the direct mediated mechanism is independent of the prior distribution F .

The modified mediator messages and the direct mediated allocation function h^d essentially transfer the randomness in the strategies of the players to the mediator messages, thus allowing each player to simply report her type. We observe that the truthful strategies

$$\tau_i^d(\tilde{\theta}_i | \phi'_i, \theta_i) = \mathbf{1}\{\tilde{\theta}_i = \theta_i\},$$

for all players i , implement the allocation choice function f for the direct mediated mechanism \mathcal{M}^d . Here is a formal proof.

Let us compute the distribution on the allocation set induced by the truthful strategy for the direct mediated mechanism. For any fixed $\theta \in \Theta$ and $\alpha \in A$, we have

$$\begin{aligned} & \sum_{\phi' \in \Phi'} D'(\phi') \sum_{\tilde{\theta} \in \Theta} \left(\prod_{i \in [n]} \tau_i^d(\tilde{\theta}_i | \phi'_i, \theta_i) \right) h^d(\alpha | \phi', \tilde{\theta}) \\ &= \sum_{\phi' \in \Phi'} D'(\phi') \sum_{\tilde{\theta} \in \Theta} \left(\prod_{i \in [n]} \mathbf{1}\{\tilde{\theta}_i = \theta_i\} \right) h^d(\alpha | \phi', \tilde{\theta}) \\ & \quad \dots \text{ because } \tau^d \text{ is a truthful strategy} \\ &= \sum_{\phi' \in \Phi'} D'(\phi') h^d(\alpha | \phi', \theta) \\ &= \sum_{\phi' \in \Phi'} D(\phi) \left(\prod_{i \in [n]} \prod_{\theta'_i \in \Theta_i} \tau_i \left(\psi_i^{\theta'_i} | \phi_i, \theta'_i \right) \right) h^d(\alpha | \phi', \theta) \\ & \quad \dots \text{ from (6.A.3)} \\ &= \sum_{\phi \in \Phi} D(\phi) \sum_{\substack{(\psi_i^{\theta'_i})_{\theta'_i \in \Theta_i, i \in [n]} \\ \in \prod_{i \in [n]} (\Psi_i)^{\Theta_i}}} \left(\prod_{i \in [n]} \prod_{\theta'_i \in \Theta_i} \tau_i \left(\psi_i^{\theta'_i} | \phi_i, \theta'_i \right) \right) h^d(\alpha | \phi', \theta) \\ & \quad \dots \text{ from (6.A.2)} \\ &= \sum_{\phi \in \Phi} D(\phi) \sum_{\substack{(\psi_i^{\theta'_i})_{\theta'_i \in \Theta_i, i \in [n]} \\ \in \prod_{i \in [n]} (\Psi_i)^{\Theta_i}}} \left(\prod_{i \in [n]} \prod_{\theta'_i \in \Theta_i} \tau_i \left(\psi_i^{\theta'_i} | \phi_i, \theta'_i \right) \right) h \left(\alpha | \phi, \left(\psi_i^{\theta'_i} \right)_{i \in [n]} \right) \end{aligned}$$

$$\begin{aligned}
 & \dots \text{ from (6.A.5)} \\
 & = \sum_{\phi \in \Phi} D(\phi) \sum_{\substack{(\psi_i^{\theta_i})_{i \in [n]} \\ \in \prod_{i \in [n]} \Psi_i}} \sum_{\substack{(\psi_i^{\theta'_i})_{\theta'_i \neq \theta_i, i \in [n]} \\ \in \prod_{i \in [n]} (\Psi_i)^{\Theta_i \setminus \theta_i}}} \left(\prod_{i \in [n]} \prod_{\theta'_i \neq \theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \right) \\
 & \quad \times \left(\prod_{i \in [n]} \tau_i(\psi_i^{\theta_i} | \phi_i, \theta_i) \right) h(\alpha | \phi, (\psi_i^{\theta_i})_{i \in [n]}) \\
 & = \sum_{\phi \in \Phi} D(\phi) \sum_{\substack{(\psi_i^{\theta_i})_{i \in [n]} \\ \in \prod_{i \in [n]} \Psi_i}} \left(\prod_{i \in [n]} \tau_i(\psi_i^{\theta_i} | \phi_i, \theta_i) \right) h(\alpha | \phi, (\psi_i^{\theta_i})_{i \in [n]}) \\
 & \quad \times \sum_{\substack{(\psi_i^{\theta'_i})_{\theta'_i \neq \theta_i, i \in [n]} \\ \in \prod_{i \in [n]} (\Psi_i)^{\Theta_i \setminus \theta_i}}} \left(\prod_{i \in [n]} \prod_{\theta'_i \neq \theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \right) \\
 & = \sum_{\phi \in \Phi} D(\phi) \sum_{\substack{(\psi_i^{\theta_i})_{i \in [n]} \\ \in \prod_{i \in [n]} \Psi_i}} \left(\prod_{i \in [n]} \tau_i(\psi_i^{\theta_i} | \phi_i, \theta_i) \right) h(\alpha | \phi, (\psi_i^{\theta_i})_{i \in [n]}) \\
 & \quad \times \left(\prod_{i \in [n]} \prod_{\theta'_i \neq \theta_i} \sum_{\psi_i^{\theta'_i} \in \Psi_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \right) \\
 & = \sum_{\phi \in \Phi} D(\phi) \sum_{\substack{(\psi_i^{\theta_i})_{i \in [n]} \\ \in \prod_{i \in [n]} \Psi_i}} \left(\prod_{i \in [n]} \tau_i(\psi_i^{\theta_i} | \phi_i, \theta_i) \right) h(\alpha | \phi, (\psi_i^{\theta_i})_{i \in [n]}) \left(\prod_{i \in [n]} \prod_{\theta'_i \neq \theta_i} 1 \right) \\
 & \quad \dots \text{ because } \tau_i(\cdot | \phi_i, \theta'_i) \in \Delta(\Psi_i) \\
 & = \sum_{\phi \in \Phi} D(\phi) \sum_{\psi \in \Psi} \left(\prod_{i \in [n]} \tau_i(\psi_i | \phi_i, \theta_i) \right) h(\alpha | \phi, \psi) \\
 & = f(\alpha | \theta) \text{ if } \theta \in \text{supp } F \\
 & \quad \dots \text{ from (6.4.8)}.
 \end{aligned}$$

This confirms that the truthful strategy profile implements the social choice function for the direct mediated mechanism \mathcal{M}^d .

We will now show that if τ is an F -Bayes-Nash equilibrium for \mathcal{M} , then τ^d is an F -Bayes-Nash equilibrium for \mathcal{M}^d . We will then show that if τ is a belief-dominant equilibrium

for \mathcal{M} , then τ^d is a belief-dominant equilibrium for \mathcal{M}^d . To prove these two statements, we first make the following observation concerning the lottery induced over the allocations for player i in the setting of the direct mediated mechanism \mathcal{M}^d , when she receives the message $\phi'_i := (\phi_i, (\psi_i^{\theta'_i})_{\theta'_i \in \Theta_i}) \in \text{supp } D'_i$, has type $\theta_i \in \Theta_i$, has a belief $G'_{-i} \in \Delta(\Theta_{-i})$ on the opponents' type reports (which are the signals of the opponents in this direct mediated mechanism), and decides to report $\tilde{\theta}_i$. The lottery induced over the allocations for player i satisfies

$$\begin{aligned} \mu'_i(\phi'_i, \theta_i, \tilde{\theta}_i; \mathcal{M}^d, G'_{-i}) &:= \sum_{\phi'_{-i} \in \Phi'_{-i}} D'_{-i}(\phi'_{-i} | \phi'_i) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h^d(\phi', \tilde{\theta}_i, \theta_{-i}) \\ &= \sum_{\phi_{-i} \in \Phi_{-i}} D_{-i}(\phi_{-i} | \phi_i) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \\ &\quad \times \sum_{\psi_{-i}} \left(\prod_{j \neq i} \tau_j(\psi_j | \phi_j, \theta_j) \right) h(\phi, \psi_i^{\tilde{\theta}_i}, \psi_{-i}). \end{aligned} \quad (6.A.6)$$

We give a formal proof of this in Appendix 6.B. Let us see how this observation helps us prove the two statements above, namely, τ^d is an equilibrium (F -Bayes-Nash or belief-dominant resp.) of \mathcal{M}^d given that τ is an equilibrium (F -Bayes-Nash or belief-dominant resp.) of \mathcal{M} .

Suppose F is the common prior and τ is an F -Bayes-Nash equilibrium for the mediated mechanism \mathcal{M} . Let $\phi'_i \in \text{supp } D'_i$ and $\theta_i \in \text{supp } F_i$. From (6.A.4), we know that $D'_i(\phi'_i) > 0$ implies $D_i(\phi_i) > 0$ and $\tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) > 0$, for all $\theta'_i \in \Theta_i$, (and in particular, we have $\tau_i(\psi_i^{\theta_i} | \phi_i, \theta_i) > 0$). Since τ is a Bayes-Nash equilibrium for \mathcal{M} , we have

$$W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi_i^{\theta_i}; \mathcal{M}, F, \tau_{-i})) \geq W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \tilde{\psi}_i; \mathcal{M}, F, \tau_{-i})),$$

for all $\tilde{\psi}_i \in \Psi_i$. (Note that $\psi_i^{\theta_i} \in \text{supp } \tau_i(\cdot | \phi_i, \theta_i)$, $\phi_i \in \text{supp } D_i$, $\theta_i \in \text{supp } F_i$.) Taking $G'_{-i} = F_{-i}(\cdot | \theta_i)$ in (6.A.6), we get that

$$\mu'_i(\phi'_i, \theta_i, \tilde{\theta}_i; \mathcal{M}^d, F, \tau_{-i}^d) = \mu_i(\phi_i, \theta_i, \psi_i^{\tilde{\theta}_i}; \mathcal{M}, F, \tau_{-i}), \quad (6.A.7)$$

for all $\tilde{\theta}_i \in \Theta_i$, and thus,

$$\begin{aligned} W_i^{\theta_i}(\mu'_i(\phi'_i, \theta_i, \tilde{\theta}_i; \mathcal{M}^d, F, \tau_{-i}^d)) &= W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi_i^{\tilde{\theta}_i}; \mathcal{M}, F, \tau_{-i})) \\ &\geq W_i^{\theta_i}(\mu_i(\phi_i, \theta_i, \psi_i^{\tilde{\theta}_i}; \mathcal{M}, F, \tau_{-i})) \\ &= W_i^{\theta_i}(\mu'_i(\phi'_i, \theta_i, \tilde{\theta}_i; \mathcal{M}^d, F, \tau_{-i}^d)), \end{aligned} \quad (6.A.8)$$

for all $\tilde{\theta}_i \in \Theta_i$. This establishes that the truthful strategy τ^d is an F -Bayes-Nash equilibrium for \mathcal{M} .

Now suppose τ is a belief-dominant strategy for \mathcal{M} . Let $\phi'_i \in \text{supp } D'_i$ and $\theta_i \in \Theta_i$. Again, this implies $D_i(\phi_i) > 0$ and $\psi_i^{\theta_i} \in \text{supp } \tau_i(\phi_i, \theta_i)$. Corresponding to a belief $G'_{-i} \in \Delta(\Theta_{-i})$, consider the belief $G_{-i} \in \Delta(\Psi_{-i})$ given by

$$G_{-i}(\psi_{-i}) := \sum_{\phi_{-i} \in \Phi_{-i}} D_{-i}(\phi_{-i} | \phi_i) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \left(\prod_{j \neq i} \tau_j(\psi_j | \phi_j, \theta_j) \right) \quad (6.A.9)$$

Then, from (6.A.6), we have that

$$\mu'_i(\phi'_i, \theta_i, \tilde{\theta}_i; \mathcal{M}^d, G'_{-i}) = \mu_i(\phi_i, \theta_i, \psi_i^{\tilde{\theta}_i}; \mathcal{M}, G_{-i}) \quad (6.A.10)$$

Noting that $\psi_i^{\tilde{\theta}_i} \in \text{supp } \tau_i(\phi_i, \tilde{\theta}_i)$ for all $\tilde{\theta}_i \in \Theta_i$ and $\phi_i \in \text{supp } D_i$, and τ_i being a belief-dominant strategy, we get that

$$W_i^{\theta_i}(\mu'_i(\phi'_i, \theta_i, \theta_i; \mathcal{M}^d, G'_{-i})) \geq W_i^{\theta_i}(\mu'_i(\phi'_i, \theta_i, \tilde{\theta}_i; \mathcal{M}^d, G'_{-i})) \quad (6.A.11)$$

for all $\tilde{\theta}_i \in \Theta_i$. Thus, the truthful strategy τ^d is a belief-dominant strategy for \mathcal{M}^d .

This completes the proof of statement (i) in Theorem 6.4.1 and part of statement (iii) corresponding to mediated mechanisms. We now consider the setting of publicly mediated mechanisms and establish the rest of the theorem.

Let

$$\mathcal{M}_* = (\Phi_*, D_*, (\Psi_i)_{i \in [n]}, h_*)$$

be a publicly mediated mechanism and for each player i let $\tau_i : \Phi_* \times \Theta_i \rightarrow \Delta(\Psi_i)$ be her strategy such that the strategy profile τ induces the allocation choice function f for this mechanism. We now consider the direct publicly mediated mechanism

$$\mathcal{M}_*^d := (\Phi_*', D_*', (\Theta_i)_{i \in [n]}, h_*^d),$$

where the message set is given by

$$\Phi_*' := \Phi_* \times \prod_{i=1}^n (\Psi_i)^{\Theta_i},$$

with a typical element denoted by

$$\phi_*' := (\phi_*, (\psi_i^{\theta'_i})_{\theta'_i \in \Theta_i, i \in [n]}), \quad (6.A.12)$$

and the mediator distribution D_*' is given by

$$D_*'(\phi_*') := D_*(\phi_*) \prod_{i \in [n]} \prod_{\theta'_i \in \Theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \text{ for all } \phi_*' \in \Phi_*'. \quad (6.A.13)$$

Similar to the previous setting, here the modified mediator messages and the mediator distribution can be interpreted as encapsulating the randomness in the strategies of the players for each of their types into the public messages. We can similarly verify that D'_* is indeed a probability distribution on Φ'_* . The direct mediated allocation function h_*^d in the direct publicly mediated mechanism \mathcal{M}^d is given by

$$h_*^d(\phi'_*, \theta') := h_* \left(\phi_*, \left(\psi_i^{\theta'_i} \right)_{i \in [n]} \right) \text{ for all } \phi'_* \in \Phi'_*, \theta' \in \Theta. \quad (6.A.14)$$

We can similarly verify that the truthful strategies

$$\tau^d(\phi'_*, \theta_i) = \theta_i$$

implement the allocation choice function f for \mathcal{M}_*^d .

Fix $\phi'_* \in \text{supp } D'_*$. Note that

$$h_*^d(\phi'_*, \tilde{\theta}_i, \theta_{-i}) = h_*(\phi_*, \psi_i^{\tilde{\theta}_i}, (\psi_j^{\theta_j})_{j \neq i}), \quad (6.A.15)$$

for all $\tilde{\theta}_i \in \Theta_i$. From (6.A.13), we have $\phi_* \in \text{supp } D_*$ and $\psi_i^{\theta_i} \in \text{supp } \tau_i(\phi_*, \theta_i)$ for all $\theta_i \in \Theta_i$.

Now suppose τ is a dominant equilibrium for \mathcal{M}_* . The lottery induced over the allocations for player i when she receives a publicly mediated message ϕ'_* , has type θ_i , believes that the opponents are reporting θ_{-i} , and decides to report $\tilde{\theta}_i$ is given by

$$\mu'_i(\phi'_*, \theta_i, \tilde{\theta}_i; \mathcal{M}_*^d, \theta_{-i}) = h_*^d(\phi'_*, \tilde{\theta}_i, \theta_{-i}). \quad (6.A.16)$$

We get this from (6.4.9) by considering the special case of publicly mediated mechanisms. From (6.A.15), we get that this is equal to the lottery induced over the allocations for player i when she receives a publicly mediated message ϕ_* , has type θ_i , believes that the opponents are reporting $\psi_j^{\theta_j}$, $j \neq i$, and decides to report $\psi_i^{\tilde{\theta}_i}$, namely,

$$\mu_i(\phi_*, \theta_i, \psi_i^{\tilde{\theta}_i}; \mathcal{M}_*, (\psi_j^{\theta_j})_{j \neq i}) = h_*(\phi_*, \psi_i^{\tilde{\theta}_i}, (\psi_j^{\theta_j})_{j \neq i}).$$

Since τ_i is a dominant strategy, $\phi_* \in \text{supp } D_*$, and $\psi_i^{\theta_i} \in \text{supp } \tau_i(\phi_*, \theta_i)$, we have

$$W_i^{\theta_i}(\mu_i(\phi_*, \theta_i, \psi_i^{\theta_i}; \mathcal{M}_*, (\psi_j^{\theta_j})_{j \neq i})) \geq W_i^{\theta_i}(\mu_i(\phi_*, \theta_i, \tilde{\psi}_i; \mathcal{M}_*, (\psi_j^{\theta_j})_{j \neq i})),$$

for all $\tilde{\psi}_i \in \Psi_i$. Hence, we have

$$W_i^{\theta_i}(\mu'_i(\phi'_*, \theta_i, \theta_i; \mathcal{M}_*^d, \theta_{-i})) \geq W_i^{\theta_i}(\mu'_i(\phi'_*, \theta_i, \tilde{\theta}_i; \mathcal{M}_*^d, \theta_{-i})),$$

for all $\tilde{\theta}_i \in \Theta_i$. Thus, τ^d is a dominant equilibrium of \mathcal{M}_*^d .

Now suppose that τ is a belief-dominant equilibrium for \mathcal{M}_* . Consider the fixed message

$$\phi'_* = (\phi_*, (\psi_i^{\theta'_i})_{\theta'_i \in \Theta_i, i \in [n]}) \in \text{supp } D'_*,$$

as before. Corresponding to a belief $G'_{-i} \in \Delta(\Theta_{-i})$, consider $G_{-i,*} \in \Delta(\Psi_{-i})$ given by

$$G_{-i,*}(\tilde{\psi}_{-i}) := \sum_{\substack{\theta_{-i} \in \Theta_{-i} \\ \text{s.t. } \psi_j^{\theta_j} = \tilde{\psi}_{-i}, \forall j \neq i}} G'_{-i}(\theta_{-i}), \quad (6.A.17)$$

for all $\tilde{\psi}_{-i} \in \Psi_{-i}$, where $\psi_j^{\theta_j}$ are the signals corresponding to the types as defined by the message ϕ'_* .

As observed in equation (6.A.16), the lottery induced over the allocations for player i , when she receives message ϕ'_* , has type θ_i , believes that the opponents' are reporting θ_{-i} , and decides to report $\tilde{\theta}_i$ is given by $h_*^d(\phi'_*, \tilde{\theta}_i, \theta_{-i})$. Now suppose that she has belief G'_{-i} on her opponents' type report instead. Then, the induced lottery over the allocations for player i is given by

$$\begin{aligned} \mu'_i(\phi'_*, \theta_i, \tilde{\theta}_i; \mathcal{M}_*^d, G'_{-i}) &= \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h_*^d(\phi'_*, \tilde{\theta}_i, \theta_{-i}) \\ &= \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h_*(\phi_*, \psi_i^{\tilde{\theta}_i}, (\psi_j^{\theta_j})_{j \neq i}) \\ &\quad \dots \text{ from (6.A.15)} \\ &= \sum_{\tilde{\psi}_{-i} \in \Psi_{-i}} h_*(\phi_*, \psi_i^{\tilde{\theta}_i}, \tilde{\psi}_{-i}) \sum_{\substack{\theta_{-i} \in \Theta_{-i} \\ \text{s.t. } \psi_j^{\theta_j} = \tilde{\psi}_{-i}, \forall j \neq i}} G'_{-i}(\theta_{-i}) \\ &= \sum_{\tilde{\psi}_{-i} \in \Psi_{-i}} h_*(\phi_*, \psi_i^{\tilde{\theta}_i}, \tilde{\psi}_{-i}) G_{-i,*}(\tilde{\psi}_{-i}) \\ &\quad \dots \text{ from (6.A.17)} \\ &= \mu_i(\phi_*, \theta_i, \psi_i^{\tilde{\theta}_i}; \mathcal{M}_*, G_{-i,*}). \end{aligned}$$

Since τ is a belief-dominant equilibrium, $\phi_* \in \text{supp } D_*$, and $\psi_i^{\theta_i} \in \text{supp } \tau_i(\phi_*, \theta_i)$ for all $\theta_i \in \Theta_i$, we have

$$W_i^{\theta_i}(\mu_i(\phi_*, \theta_i, \psi_i^{\theta_i}; \mathcal{M}_*, G_{-i,*}) \geq W_i^{\theta_i}(\mu_i(\phi_*, \theta_i, \tilde{\psi}_i; \mathcal{M}_*, G_{-i,*}),$$

for all $\tilde{\psi}_i \in \Psi_i$. Hence, we have

$$W_i^{\theta_i}(\mu'_i(\phi'_*, \theta_i, \theta_i; \mathcal{M}_*^d, G'_{-i})) \geq W_i^{\theta_i}(\mu'_i(\phi'_*, \theta_i, \tilde{\theta}_i; \mathcal{M}_*^d, G'_{-i})),$$

for all $\tilde{\theta}_i \in \Theta_i$. Thus, τ^d is a belief-dominant equilibrium of \mathcal{M}_*^d .

This completes the proof of the theorem.

6.B Proof of Equation (6.A.6)

Let us recall the first setting considered in Appendix 6.A. We have a mediated mechanism

$$\mathcal{M} = ((\Phi_i)_{i \in [n]}, D, (\Psi_i)_{i \in [n]}, h),$$

and a corresponding strategy profile τ . We had constructed a direct mediated mechanism

$$\mathcal{M}^d = ((\Phi'_i)_{i \in [n]}, D', (\Theta_i)_{i \in [n]}, h^d),$$

given by (6.A.1), (6.A.2), (6.A.3), and (6.A.5). We are interested in the situation when player i receives message $\phi'_i := (\phi_i, (\psi_i^{\theta'_i})_{\theta'_i \in \Theta_i}) \in \text{supp } D'_i$, has type $\theta_i \in \Theta_i$, and belief $G'_{-i} \in \Delta(\Theta_{-i})$ on the opponents' type reports, and decides to report $\tilde{\theta}_i$. Since $D'(\phi'_i) > 0$ by assumption, we have

$$\begin{aligned} & \sum_{\phi'_{-i} \in \Phi'_{-i}} D'_{-i}(\phi'_{-i} | \phi'_i) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h^d(\phi', \tilde{\theta}_i, \theta_{-i}) \\ &= \frac{\sum_{\phi'_{-i} \in \Phi'_{-i}} D'(\phi'_i, \phi'_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h^d(\phi', \tilde{\theta}_i, \theta_{-i})}{\sum_{\phi'_{-i} \in \Phi'_{-i}} D'(\phi'_i, \phi'_{-i})}. \end{aligned}$$

Let the denominator be denoted by

$$C_1 := \sum_{\phi'_{-i} \in \Phi'_{-i}} D'(\phi'_i, \phi'_{-i}) = D'_i(\phi'_i).$$

We now focus on the numerator, to get

$$\begin{aligned} & \sum_{\phi'_{-i} \in \Phi'_{-i}} D'(\phi'_i, \phi'_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h^d(\phi', \tilde{\theta}_i, \theta_{-i}) \\ &= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\substack{(\psi_j^{\theta'_j})_{\theta'_j \in \Theta_j, j \neq i} \\ \in \prod_{j \neq i} (\Psi_j)^{\Theta_j}}} \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j} \tau_j(\psi_j^{\theta'_j} | \phi_j, \theta'_j) \right) \prod_{\theta'_i \in \Theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \\ & \quad \times \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h^d(\phi', \tilde{\theta}_i, \theta_{-i}) \\ & \quad \dots \text{ from (6.A.3)} \\ &= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\substack{(\psi_j^{\theta'_j})_{\theta'_j \in \Theta_j, j \neq i} \\ \in \prod_{j \neq i} (\Psi_j)^{\Theta_j}}} \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j} \tau_j(\psi_j^{\theta'_j} | \phi_j, \theta'_j) \right) \left(\prod_{\theta'_i \in \Theta_i} \tau_i(\psi_i^{\theta'_i} | \phi_i, \theta'_i) \right) \end{aligned}$$

$$\begin{aligned}
& \times \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h \left(\phi_i, \phi_{-i}, \psi_i^{\bar{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right) \\
& \hspace{15em} \dots \text{ from (6.A.5)} \\
= & \left(\prod_{\theta'_i \in \Theta_i} \tau_i \left(\psi_i^{\theta'_i} | \phi_i, \theta'_i \right) \right) \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\substack{(\psi_j^{\theta'_j})_{\theta'_j \in \Theta_j, j \neq i} \\ \in \prod_{j \neq i} (\Psi_j)^{\Theta_j}}} \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j} \tau_j \left(\psi_j^{\theta'_j} | \phi_j, \theta'_j \right) \right) \\
& \times \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h \left(\phi_i, \phi_{-i}, \psi_i^{\bar{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right)
\end{aligned}$$

Let

$$C_2 := \prod_{\theta'_i \in \Theta_i} \tau_i \left(\psi_i^{\theta'_i} | \phi_i, \theta'_i \right).$$

We have,

$$\begin{aligned}
& \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\substack{(\psi_j^{\theta'_j})_{\theta'_j \in \Theta_j, j \neq i} \\ \in \prod_{j \neq i} (\Psi_j)^{\Theta_j}}} \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j} \tau_j \left(\psi_j^{\theta'_j} | \phi_j, \theta'_j \right) \right) \\
& \hspace{15em} \times \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) h \left(\phi_i, \phi_{-i}, \psi_i^{\bar{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right) \\
= & \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \sum_{\substack{(\psi_j^{\theta'_j})_{\theta'_j \in \Theta_j, j \neq i} \\ \in \prod_{j \neq i} (\Psi_j)^{\Theta_j}}} \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j} \tau_j \left(\psi_j^{\theta'_j} | \phi_j, \theta'_j \right) \right) \\
& \hspace{15em} \times h \left(\phi_i, \phi_{-i}, \psi_i^{\bar{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right) \\
= & \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \\
& \times \sum_{\substack{(\psi_j^{\theta'_j})_{j \neq i} \\ \in (\Psi_j)_{j \neq i}}} \sum_{\substack{(\psi_j^{\theta'_j})_{\theta'_j \in \Theta_j \setminus \theta_j, j \neq i} \\ \in \prod_{j \neq i} (\Psi_j)^{\Theta_j \setminus \theta_j}}} \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j \setminus \theta_j} \tau_j \left(\psi_j^{\theta'_j} | \phi_j, \theta'_j \right) \right) \\
& \times \left(\prod_{j \neq i} \tau_j \left(\psi_j^{\theta_j} | \phi_j, \theta_j \right) \right) h \left(\phi_i, \phi_{-i}, \psi_i^{\bar{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right)
\end{aligned}$$

$$\begin{aligned}
&= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \\
&\quad \times \sum_{\substack{(\psi_j^{\theta_j})_{j \neq i} \\ \in (\Psi_j)_{j \neq i}}} \left(\prod_{j \neq i} \tau_j \left(\psi_j^{\theta_j} | \phi_j, \theta_j \right) \right) h \left(\phi_i, \phi_{-i}, \psi_i^{\tilde{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right) \\
&\quad \times \sum_{\substack{(\psi_j^{\theta'_j})_{\theta'_j \in \Theta_j \setminus \theta_j, j \neq i} \\ \in \prod_{j \neq i} (\Psi_j)^{\Theta_j \setminus \theta_j}}} \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j \setminus \theta_j} \tau_j \left(\psi_j^{\theta'_j} | \phi_j, \theta'_j \right) \right) \\
&= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \\
&\quad \times \sum_{\substack{(\psi_j^{\theta_j})_{j \neq i} \\ \in (\Psi_j)_{j \neq i}}} \left(\prod_{j \neq i} \tau_j \left(\psi_j^{\theta_j} | \phi_j, \theta_j \right) \right) h \left(\phi_i, \phi_{-i}, \psi_i^{\tilde{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right) \\
&\quad \times \left(\prod_{j \neq i} \prod_{\theta'_j \in \Theta_j \setminus \theta_j} \sum_{\psi_j^{\theta'_j} \in \Psi_j} \tau_j \left(\psi_j^{\theta'_j} | \phi_j, \theta'_j \right) \right) \\
&= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \\
&\quad \times \sum_{\substack{(\psi_j^{\theta_j})_{j \neq i} \\ \in (\Psi_j)_{j \neq i}}} \left(\prod_{j \neq i} \tau_j \left(\psi_j^{\theta_j} | \phi_j, \theta_j \right) \right) h \left(\phi_i, \phi_{-i}, \psi_i^{\tilde{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right) \\
&\quad \times \left(\prod_{j \neq i} \prod_{\tilde{\theta}'_j \in \Theta_j \setminus \theta_j} 1 \right) \\
&= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \\
&\quad \times \sum_{\substack{(\psi_j^{\theta_j})_{j \neq i} \\ \in (\Psi_j)_{j \neq i}}} \left(\prod_{j \neq i} \tau_j \left(\psi_j^{\theta_j} | \phi_j, \theta_j \right) \right) h \left(\phi_i, \phi_{-i}, \psi_i^{\tilde{\theta}_i}, \left(\psi_j^{\theta_j} \right)_{j \neq i} \right) \\
&= \sum_{\phi_{-i} \in \Phi_{-i}} D(\phi_i, \phi_{-i}) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i})
\end{aligned}$$

$$\begin{aligned}
& \times \sum_{\psi_{-i}} \left(\prod_{j \neq i} \tau_j(\psi_j | \phi_j, \theta_j) \right) h(\phi_i, \phi_{-i}, \psi_i^{\tilde{\theta}_i}, \psi_{-i}) \\
= & D_i(\phi_i) \sum_{\phi_{-i} \in \Phi_{-i}} D_{-i}(\phi_{-i} | \phi_i) \sum_{\theta_{-i} \in \Theta_{-i}} G'_{-i}(\theta_{-i}) \\
& \times \sum_{\psi_{-i}} \left(\prod_{j \neq i} \tau_j(\psi_j | \phi_j, \theta_j) \right) h(\phi_i, \phi_{-i}, \psi_i^{\tilde{\theta}_i}, \psi_{-i}).
\end{aligned}$$

We recall that $D_i(\phi_i)C_2/C_1 = 1$ from (6.A.4), and hence, we get (6.A.6).

6.C Outcome Sets can be Identified with the Allocation Set under EUT

Consider a setting in which all the players have EUT preferences for all their types. For this restricted setting, we will now construct an environment

$$\mathcal{E}' := ([n], (\Theta'_i)_{i \in [n]}, A, (\Gamma'_i)_{i \in [n]}, \zeta'),$$

that we call the reduced environment corresponding to the environment (as defined in (6.2.1))

$$\mathcal{E} := ([n], (\Theta_i)_{i \in [n]}, A, (\Gamma_i)_{i \in [n]}, \zeta).$$

From (6.2.9), we observe that, since we are dealing with EUT preferences, the utility function $W_i^{\theta_i}$ is completely determined by the values $u_i^{\theta_i}(\alpha), \forall \alpha \in A$. Suppose the mechanism designer models the outcome set of each player i by $\Gamma'_i = A$ instead of the true outcome set Γ_i , with the trivial allocation-outcome mapping ζ'_i instead of the original allocation-outcome mapping ζ_i . Let ζ' denote the product of the trivial allocation-outcome mappings $\zeta'_i, i \in [n]$. Corresponding to a type $\theta_i \in \Theta_i$ for player i , the mechanism designer models her type by θ'_i , which is characterized by the utility function $u_i^{\theta'_i} : \Gamma'_i \rightarrow \mathbb{R}$ as defined in (6.2.8). Since the players are assumed to have EUT preferences, the probability weighting functions under each type θ'_i are modeled to be $w_i^{\pm}(p) = p, \forall p \in [0, 1]$. Let Θ'_i denote the set comprised of all types θ'_i corresponding to the types $\theta_i \in \Theta_i$. Let $\mathcal{T}_i : \Theta_i \rightarrow \Theta'_i$ denote the function for this correspondence. Suppose the mechanism designer treats the environment as if given by

$$\mathcal{E}' := ([n], (\Theta'_i)_{i \in [n]}, A, (\Gamma'_i)_{i \in [n]}, \zeta').$$

Let $\Theta' := \prod_i \Theta'_i$. Let $\mathcal{T} : \Theta \rightarrow \Theta'$ denote the product transformation defined by the functions $\mathcal{T}_i, i \in [n]$. Notice that the function \mathcal{T}_i is a bijection since, as pointed out earlier, even if $u_i^{\theta_i} = u_i^{\tilde{\theta}_i}$ for some $\theta_i \neq \tilde{\theta}_i$, we will treat $\mathcal{T}_i(\theta_i)$ and $\mathcal{T}_i(\tilde{\theta}_i)$ as different elements of Θ'_i . For any prior $F \in \Delta(\Theta)$, let $F' \in \Delta(\Theta')$ be the corresponding prior induced by the bijection \mathcal{T} .

Note that, for any player i , having any type θ_i , and any lottery $\mu \in \Delta(A)$, we have

$$W_i^{\theta_i}(\mu) = W_i^{\mathcal{T}_i(\theta_i)}(\mu). \quad (6.C.1)$$

(Here, $W_i^{\mathcal{T}_i(\theta_i)}$ should be interpreted as the utility function for player i with type $\theta'_i = \mathcal{T}_i(\theta_i)$ corresponding to the reduced environment \mathcal{E}' .) Let $f' : \Theta' \rightarrow A$ be an allocation choice function that is implementable in F' -Bayes-Nash equilibrium $\sigma' := (\sigma'_i)_{i \in [n]}$ (where $\sigma'_i : \Theta'_i \rightarrow \Delta(\Psi'_i)$) for the mechanism

$$\mathcal{M}_0 = ((\Psi_i)_{i \in [n]}, h_0).$$

Now suppose the system operator uses the same mechanism \mathcal{M}_0 in environment \mathcal{E} . Consider the allocation choice function $f : \Theta \rightarrow \Delta(A)$ given by

$$f(\theta) = f'(\mathcal{T}(\theta)).$$

For each player i , consider the strategy $\sigma_i : \Theta_i \rightarrow \Delta(\Psi_i)$ given by

$$\sigma_i(\theta_i) = \sigma'_i(\mathcal{T}_i(\theta_i)).$$

Similar to (6.2.13), for any $\theta'_i \in \text{supp } F'_i$ and signal ψ_i , let

$$\mu'_i(\theta'_i, \psi_i; \mathcal{M}_0, F', \sigma'_{-i}) := \sum_{\theta_{-i} \in \Theta_{-i}} F'_{-i}(\theta_{-i} | \theta_i) \sum_{\psi_{-i} \in \Psi_{-i}} \prod_{j \neq i} \sigma_j(\psi_j | \theta'_j) h_0(\psi), \quad (6.C.2)$$

be the belief of player i on the allocation set corresponding to the reduced environment \mathcal{E}' . Note that

$$\mu_i(\theta_i, \psi_i; \mathcal{M}_0, F, \sigma_{-i}) = \mu'_i(\mathcal{T}_i(\theta_i), \psi_i; \mathcal{M}_0, F', \sigma'_{-i}).$$

From observation (6.C.1) and the definition of F -Bayes-Nash equilibrium in (6.2.14) and (6.2.15), we get that the allocation choice function f is implementable in F -Bayes-Nash equilibrium by the mechanism \mathcal{M}_0 with the equilibrium strategy σ .

On the other hand, suppose we have an allocation choice function $f : \Theta \rightarrow \Delta(A)$. Consider the corresponding allocation choice function $f' : \Theta' \rightarrow \Delta(A)$ given by

$$f'(\theta') = f(\mathcal{T}^{-1}(\theta')).$$

We now observe that if f is implementable in F -Bayes-Nash equilibrium by a mechanism \mathcal{M}_0 and an F -Bayes-Nash equilibrium σ , then so is f' by the same mechanism \mathcal{M}_0 and the F' -Bayes-Nash equilibrium σ' comprised of

$$\sigma'_i(\theta'_i) = \sigma_i(\mathcal{T}_i^{-1}(\theta'_i)),$$

for all $i \in [n]$, $\theta'_i \in \Theta'_i$.

We can similarly show that if f' is implementable in dominant (resp. belief-dominant) equilibrium by a mechanism \mathcal{M}_0 with the equilibrium strategy profile σ' for the reduced environment \mathcal{E}' , then so is f by the same mechanism \mathcal{M}_0 with the corresponding equilibrium strategy profile σ for the environment \mathcal{E} , and vice versa.

Hence, under EUT, from the mechanism designer's point of view, it is enough to model the types of player i by setting the outcome set $\Gamma'_i = A$, assuming the trivial allocation-outcome mapping ζ'_i , and the types $\theta'_i \in \Theta'_i$.

Notes

¹⁸ Myerson [93] refers to the mechanism design framework as a generalized principal-agent problem. In contrast to our framework, Myerson is interested in problems where the agents have private decision domains in addition to private information. Here, by private decision domains, we mean possible actions for the player that directly affect the outcomes. For example, in employment contracts the actions of the employee directly affect the outcome. These actions should not be confused with the signals of the player in the communication protocol set up by the system operator. We prefer to call the entity in control as the system operator instead of the principal to emphasize that the system operator alone controls the system implementation. We restrict ourselves to situations where agents do not have private decision domains because such situations involve dynamic decision-making, and non-EUT models face several issues in such situations (see Section ?? for more on this). Thus our model cannot account for moral hazard.

¹⁹ Here, strictly speaking, given a matching by the platform, the users can refuse to go through with the matching. Although these decisions fall under the separate decision domains of the agents, they are rare and can be accounted for separately.

²⁰ Since we have assumed that the type of a player completely determines her CPT features, we are implicitly assuming *private preferences*, i.e. the preference over lotteries on the outcome set for each player is her private information and does not depend on other players' information or types, also known as *informational externalities* (see [134]).

²¹ Even if $u_i^{\theta_i} = u_i^{\tilde{\theta}_i}$ for some $\theta_i \neq \tilde{\theta}_i$ it is sometimes convenient to retain the connection to the underlying type. Notice that we have allowed different types of player i to have the same CPT features. Later, when we discuss mechanism design with a common prior, which is a distribution on the types of all the players, it will let us differentiate between the types of players that have identical CPT features but distinct beliefs on the opponents' types. Mechanism design often focuses on “naive type sets”, that is, the type set Θ_i for each player i is assumed to be comprised of exactly one element for each “preference type” of the player. Here, by preference type of a player we mean the preferences of the player on her outcome set. We borrow the expression “naive type sets” from [20]. In this chapter, we do not assume the type sets to be naive. Such an assumption would entail a bijective correspondence between the types θ_i and the CPT features (v_i, w_i^\pm) for each player i . This distinction is relevant because besides having a preference type, a player can also have a “belief type”. For example, the prior F could be such that $F_{-i}(\theta_i) \neq F_{-i}(\tilde{\theta}_i)$ even when the value function and the probability weighting functions corresponding to the types θ_i and $\tilde{\theta}_i$ coincide. (For more on this, see [15, 80], and Chapter 10 of [19].)

²² Note that, in general, the preferences defined by the utility function W_i^θ over the lotteries over the allocation set may not be given by CPT preferences directly, i.e. there need not exist any probability weighting functions \tilde{w}_i^\pm such that, for all $\mu \in \Delta(A)$, $W_i^{\theta_i}(\mu)$ is equal to the CPT value corresponding to the value function $u_i^{\theta_i}$ on A and the probability weighting functions \tilde{w}_i^\pm . To see this, consider a type θ_i for player i such that $V_i^{\theta_i}(L'_i) = V_i^{\theta_i}(L''_i) > V_i^{\theta_i}(0.5L'_i + 0.5L''_i)$, for lotteries $L'_i, L''_i \in \Delta(\Gamma_i)$. See [106] for an example of CPT preferences and lotteries (over 4 outcomes) that satisfy the above condition. Let there be two allocations α' and α'' such that $\zeta_i(\alpha') = L'_i$ and $\zeta_i(\alpha'') = L''_i$. If $W_i^{\theta_i}$ were to correspond to any CPT preference directly on the allocation set then, by the first order stochastic dominance property of CPT, we would get $W_i^{\theta_i}(0.5\alpha' + 0.5\alpha'') = W_i^{\theta_i}(\alpha') = W_i^{\theta_i}(\alpha'')$. But, since this is not true for the setting under consideration, we get that $W_i^{\theta_i}$ cannot correspond to any CPT preference directly over A .

²³ This is the version of the revelation principle commonly referred to in the mechanism design context. Another version of the revelation principle appears in the context of correlated equilibrium [6, 5]. This is concerned with an n -player non-cooperative game in normal form. A mediator draws a message profile, comprised of a message for each player, from a fixed joint probability distribution on the set of message profiles, and sends each player her corresponding message. The joint distribution over message profiles used is assumed to be common knowledge between the mediator and all the players. Based on her received signal, each player chooses her action (possibly from a probability distribution over her action set). When the message set for each player is the same as her action set and the probability distribution on the set of

message profiles (or equivalently action profiles) is such that truthful strategy, i.e. the strategy of choosing the action that is received as a message from the mediator, is a Nash equilibrium, then such a probability distribution is said to be a correlated equilibrium. Under EUT, the set of all correlated equilibria of a game is characterized as the union over all possible message sets and mediator distributions, of the sets of joint distributions on the action profiles of all players, arising from all the Nash equilibria for the resulting game. See [106] for a discussion on the revelation principle for correlated equilibrium when players have CPT preferences. Myerson [91] has considered a further generalization to games with incomplete information in which each player first reports her type. Analyzing such settings under CPT would entail dynamic decision making and is beyond the scope of this chapter.

Chapter 7

Concluding Remarks and Directions for Future Work

Civilization advances by
extending the number of
important operations which we
can perform without thinking
about them.

Alfred Whitehead

7.1 Introduction

EUT has been the backbone of almost every development in economics and game theory. It feels natural to extend these achievements to more general and better models of behavior if possible. CPT is an excellent tool to make progress in this direction. In my experience, CPT often introduces several modeling complications, however, in most cases, as is evident from this thesis, it is possible to extend the results from game theory and economics to players with CPT preferences. In each instance, we had to make special provisions either to the underlying framework or define appropriate notions - new ones or old ones with modified interpretations. This portrays the richness of this pursuit and a cue for future undertakings along these lines.

By no means this means that CPT can solve all the problems observed due to behavioral factors and deviations from EUT. For example, in [102], Nwogugu points out that “human beings and human decision making are subject to emotions, fairness considerations, ethical considerations, implied or actual constraints, personal aspirations, philosophical differences, regret, regret aversion, social pressure, peer pressure, phobias, perception of incentives, differences in cognition, biological differences in neural activity, willingness to defer, reciprocation, and willingness to use risk management tools, all of which result in significant departures

from rationality and traditional models of humans in decision making.” Some of these aspects can be accounted for by CPT and some cannot. This doesn’t mean that we should abandon the approach based on theories of rationality. Indeed, principled approaches based on fundamental theoretical developments have often helped in guiding real-world applications. Here is an excerpt from [118] stated in another context that illustrates the value of theoretical pursuits in applied fields:

Consider the design of suspension bridges. The Newtonian physics they embody is beautiful both in mathematics and in steel, and college students can be taught to derive the curves that describe the shape of the supporting cables. But no bridge could be built based only on this elegant theoretical treatment, in which the only force is gravity, and all beams are perfectly rigid. Real bridges are built of steel and rest on rock and soil and water, and so bridge design also concerns metal fatigue, soil mechanics, and the forces of waves and wind. Many design questions concerning these real-world complications cannot be answered analytically but, instead, must be explored using physical or computational models. Often these involve estimating magnitudes of phenomena missing from the simple Newtonian model, some of which are small enough to be of little consequence, while others will cause the bridge to fall down if not adequately addressed. Just as no suspension bridges could be built without an understanding of the underlying physics, neither could any be built without understanding many additional features, also physical in nature, but more varied and complex than addressed by the simple model. These additional features, and how they are related to and interact with that part of the physics captured by the simple model, are the concern of the scientific literature of engineering. Some of this is less elegant than the Newtonian model, but it is what makes bridges stand. Just as important, it allows bridges designed on the same basic Newtonian model to be built longer, stronger, and lighter over time, as the complexities and how to deal with them become better understood.

Roth and Peranson

“The redesign of the matching market for American physicians: Some engineering aspects of economic design.” In: *American economic review* (1999).

Such an engineering approach is essential in behavioral economics too. Today, behavioral economics is often seen as a magician’s tool used to manipulate human choices and responses in contrast to the mainstream developments in economics and game theory. For example, one of the most frequently cited examples of a nudge is the etching of the image of a housefly into the men’s room urinals at Amsterdam’s Schiphol Airport, which is intended to “improve the aim.” Some of the behavioral economists have often expressed their discomfort in this approach. For example, David Gal, a professor of marketing at the University of Illinois at Chicago, says in an article that appeared in the *New York Times* with the title “Selling

Behavioral Economics”: “There is nothing wrong with achieving small victories with minor interventions. The worry, however, is that the perceived simplicity and efficacy of such tactics will distract decision makers from more substantive efforts—for example, reducing electricity consumption by taxing it more heavily or investing in renewable energy resources. It is great that behavioral economics is receiving its due; the field has contributed significantly to our understanding of ourselves. But in all the excitement, its important to keep an eye on its limits.” George Loewenstein, a professor of economics and psychology at Carnegie Mellon University and Peter Ubel, a professor of business and public policy at Duke and the author of “Free Market Madness: Why Human Nature Is at Odds With Economics,” say in an article that appeared in the *New York Times* with the title “Economics Behaving Badly”: “Behavioral economics should complement, not substitute for, more substantive economic interventions. If traditional economics suggests that we should have a larger price difference between sugar-free and sugared drinks, behavioral economics could suggest whether consumers would respond better to a subsidy on unsweetened drinks or a tax on sugary drinks. But that’s the most it can do. For all of its insights, behavioral economics alone is not a viable alternative to the kinds of far-reaching policies we need to tackle our nations challenges.”

In my view, behavioral economics need not be limited as an addendum, but instead we must try to bring behavioral economics to the same level of rigor as classical economics. In this work, we have just scratched the surface in this regard. Already we saw several benefits of using CPT to model the players’ behavior. For example, in Chapter 2 we saw the benefits of lotteries in resource allocation which cannot be explained by EUT. Furthermore, the fact that optimal resource allocation can be achieved in a market-based setting with real-time signals between the players and the system operator would enable its implementation in real-world scenarios. In Chapter 6, we saw the use of the messaging stage to recover the revelation principle. In this chapter, we will discuss few directions for future work where behavioral economics would play a huge role if applied in the spirit of the engineering approach mentioned above. Besides, to include behavioral features not explained by CPT, we will have to incorporate other behavioral theories (some of which exist today, and some which will be developed in the future). It is important that these behavioral theories have a nice mathematical formulation. By a nice mathematical formulation, I mean something that can be used to model, study, and develop applications and algorithms for social systems.

Until now in this thesis, I have mainly restricted to making concrete statements about abstract theoretical ideas. In this chapter, at the expense of making half-baked proposals or impractical claims, I will attempt to unwrap some of the theoretical insights developed in this thesis, and provide a version of how they could play out in real-world applications. Finally, to express my intention towards this thesis and the spirit in which this work and the work to follow should be interpreted, I would like to quote Professor Rummel from his 1975 book “*Understanding Conflict and War, Vol. 1: The Dynamic Psychological Field,*”

I offer a word about my overall orientation toward this effort. I believe that to know ourselves we must focus on ourselves as individuals and in society, not

on our concrete environment, physical nature, or objective vehicles. The proper study is of our meanings, values, motives, perceptions, inner complexes, and powers. But we can know ourselves only through a particular perspective, a point of view. Whether this perspective is True, we cannot know.

However, whether my field view is the proper perspective, whether my efforts have born fruit, whether I have deluded myself about the importance of what I now have to say, should be a matter for discussion, critical evaluation, and debate. What I am offering is not Truth or the Way, but my contributions to the Struggle of Ideas out of which a better future may be forged. Our knowledge and our ability to handle our problems progress through the open conflict of ideas, through the tests of phenomenological adequacy, inner consistency, and practical-moral consequences. Reason may but err, but it can be moral. If we must err, let it be on the side of our creativity, our freedom, our betterment.

Rudolf Joseph Rummel

“Understanding Conflict and War, Vol. 1: The Dynamic Psychological Field.”
(1975)

7.2 Role of Communication, Data Analytics and Artificial Intelligence in Resource Allocation

The network resource allocation problem considered in Chapter 2 and the mechanism design framework discussed in Chapter 6 have several features in common. They comprise of a set of players which have preferences over their outcomes known privately to them. In the network resource allocation problem, the outcome for a player is the amount of resource allocated to that player, for example, bandwidth over the Internet. Resource allocation more generally would include things such as transportation units in vehicle routing or delivery systems, servers in computation networks, advertising space or visibility in social networks, or contract provisioning in financial networks or labor markets. On the other hand, player outcomes could take various forms as the delay experienced by a driver or a customer receiving a delivery, the quality of the goods or services provided to the users, or it could be the financial gains or prospects associated with the outcome. The system operator is primarily responsible for allocating resources. In the network resource allocation problem, the system operator is assumed to allocate resources to each player provided they satisfy the capacity constraints. But more generally, as considered in the mechanism design framework (Chapter 6), it implements an allocation from the set of available allocations and the implemented allocation in turn influences the outcomes for each player.

The system operator is central to this setup to facilitate optimal resource allocation. It communicates with each player, sending messages that provide important system related information to the player which affect the players beliefs and actions. These messages could

take the form of providing available options to the players, their corresponding prices, and the uncertainty associated with different outcomes. For example, a ride hailing platform provides the riders with different traveling options with estimates about their delays, service experience, and associated prices. The players respond to these messages through appropriate signaling channels provided by the system operator. These responses are governed by the players' private information about their types and their surroundings. For example, in the ride hailing example, their flexibility with time of arrival or departure, their budget, their knowledge about the people around them seeking transportation services, etc. Their responses are also affected by the behavioral traits displayed by them and the players around them and as well as their strategic policies. The system operator aggregates all these signals from the players and then allocates resources accordingly. Additionally, the system operator also maintains information about the environment and in conjunction with the information collected from the players, it is able to allocate resource more efficiently. It is evident that for the proper functioning of a system under this setup is closely tied to the communication protocols provided by the system operator and its ability to learn from them.

Notice that we are focusing our attention over markets where there is a central planner (or what we are calling a system operator) who is deciding the underlying mechanics of the system within the physical constraints and providing communication channels to the players to indicate their individual needs and preferences. Although the communication protocol often leads to a decentralized market-based mechanism, the freedom of choice for the users is restricted to the options provided by the system operator, and it is important that we maintain caution in designing these protocols. Having a central planner definitely brings the perks of improved efficiency due to potentially better planning opportunities but this is conditioned on the availability of enough relevant information. Hayek argues:

If we can agree that the economic problem of society is mainly one of rapid adaptation to changes in the particular circumstances of time and place, it would seem to follow that the ultimate decisions must be left to the people who are familiar with these circumstances, who know directly of the relevant changes and of the resources immediately available to meet them. We cannot expect that this problem will be solved by first communicating all this knowledge to a central board which, after integrating all knowledge, issues its orders. We must solve it by some form of decentralization. But this answers only part of our problem. We need decentralization because only thus can we ensure that the knowledge of the particular circumstances of time and place will be promptly used. But the "man on the spot" cannot decide solely on the basis of his limited but intimate knowledge of the facts of his immediate surroundings. There still remains the problem of communicating to him such further information as he needs to fit his decisions into the whole pattern of changes of the larger economic system.

Friedrich Hayek
"The Use of Knowledge in Society." (1945).

In the years following the period Hayek made this remark, communication technologies have made huge progress. Data collection and signaling delays have gotten reduced to microseconds with the advent of the Internet, and indeed we are seeing a lot of centralized markets in the form of big online marketplaces. But there are still several bottlenecks showing up in the form of learning useful information from all the signals and data. Besides, incorporating behavioral features in the design of these systems in a principled manner can go a long way. For example, we saw in Chapter 6, that a mediated mechanism which includes a messaging stage from the system provider to the players can help implement social choice functions which are not implementable otherwise (plus implement them truthfully by direct mechanisms). More practically, this would transform into artificial indicators from the system provider to the players that help align the players beliefs. For example, consider a food delivery platform that allows customers to order food from restaurants through their app. The platform can initiate a program where they randomly select some users and provide them with discounted service for long distance orders. As the players use this app repeatedly, they start modifying their strategies in response to this program. For example, a user would restrict his orders to local restaurants and order from far away restaurants only when he gets the message that he is among the lucky chosen customers. This is an example of a nudge provided by the platform to alter customer behavior. It is possible that such strategies would naturally come out of our mechanism design framework. Besides, our framework will help answer the question of exactly for what purpose are these nudges or incentive programs being used - How do they affect the welfare of customers? How do they affect the revenue of the platform?

When applied to real-world scenarios, it is very likely that the solution coming out from theory would require a complex signaling scheme between the system operator and the players. Naturally, human players would not be able to maneuver if such complex signaling schemes are implemented in practice. Additionally, many times the system operator does not have access to all the necessary information related to resource availability and implementation. We will now discuss ways in which these communication protocols and information collection activities can be implemented in practice. Developments in Artificial Intelligence (AI) would play a key role in these aspects.

Today, big data analytics is a hot topic that has found applications in several domains such as manufacturing, commerce, healthcare, financial services, safety and security. It is being used to:

- Predict equipment failure: Machine data such as its year of manufacturing, make, model, log entries, sensor data, error messages, engine temperature, and other factors can be used to deploy maintenance more efficiently and predict the remaining optimal life and state of systems and components.
- Assess resource availability: In situations where it is hard to get direct access to resource data, information from other sources such as user feedback can be analyzed to access this information.

- Anticipate customer demand: Data from focus groups, social media, and customer feedback, which comes in varying formats, can be used for product development, resource deployment, and operation fine-tuning to improve customer experience.
- Identify high-value customers: Insights from customer choices and spending patterns can be used to identify types of customers and use this information to target marketing strategies accordingly.
- Optimize merchandising: Analyzing data from mobile apps, in-store purchases, and geolocations will help improve inventory management and consequently encourage customers to complete purchases.
- Perform pricing analytics: Transaction data and information about supply and demand will help improve pricing strategies.
- Provide personalized recommendations: Data collected from repeated interactions with the customers can be used to provide offers that are fine-tuned to their requirements.
- Detect irregularity: Data from past behavior patterns can be used to detect fraud by identifying irregular transactions, to avoid accidents by detecting irregular driving patterns, or to caution customers against decisions they might consider irresponsible if they were in a different emotional state.

Notice how most of these tasks can be conveniently stated in our framework based on game theory, economics and behavioral psychology. Indeed, predicting equipment failure and assessing resource availability are related to the system operator gathering information about the environment such as capacity constraints in the model discussed in Chapter 2 or the allocation set and the mappings from allocation to outcomes in the mechanism design framework discussed in Chapter 6. Optimizing merchandising is a related task where the system operator actively influences the resource availability. Anticipating customer demand and identifying high-value customers relates to learning the type of players. Personalized recommendations and pricing are a part of the communication protocol between the players and the system operator. Detecting irregularity and fraud are a by-product of our behavioral approach that would help improve safety and security.

The learning tasks above such as gathering information about the environment or the players behavior and needs will involve taking advantage of the huge data collected through repeated interactions between the players and the systems, data coming from sensors and other unstructured sources like natural language or images. AI techniques such as Machine Learning (ML) and Reinforcement Learning (RL) algorithms aim to solve these problems. These algorithms require the designer to provide an objective function to maximize or a loss function to minimize. Also the communication framework is often assumed fixed exogenously either in an ad hoc fashion or relying on the designer's experience. Our holistic approach will not only guide the design of these objective functions and communication protocols but it will also incorporate the network effects coming from strategic interaction between

the players that are often missing from AI studies. Rarely is it true that the decisions and policies matter to a single individual without affecting other players in the system. Our framework will allow us to approach these problems in a principled manner and give rise to end-to-end solutions that are interpretable, efficient, and robust. Ideally, we want AI to be an accessory that would help implement the market-based strategy coming out of our framework in a practical manner by circumventing the complexity in the proposed solution. This will necessarily give rise to approximate solutions and methods from approximation theory, computation theory, and complexity theory that have gained prominence in computer science will of special interest.

7.3 Fairness and Ethical Considerations

Fairness in and of itself isn't, and shouldn't be, the goal toward which we strive. It's simply the most obvious result of a far more complex interplay of needs and systems.

Vox magazine - Fairness - May 2021

In Chapter 2, our system problem based on Kelly's work follows a utilitarian approach where the objective function is designed to maximize the total social welfare. (Another example related to the framework of mechanism design discussed in Chapter 6 is the Vickrey-Clarke-Groves (VCG) mechanism, which aims to implement a social choice function that maximizes the social welfare for each type profile.) CPT value allows us to incorporate the perceived happiness or satisfaction of the individuals when faced with uncertain prospects. In contrast to the EUT based model, it is able to capture psychological factors such as the lure of winning or the fear of losing. In addition, the reference point dependence property of CPT value provides a way to capture the behavioral expectations of the individuals such as social norms in their group or neighborhood, which is an important consideration in distributive justice. These behavioral features have significant implications to aspects such as fairness and social welfare in resource allocations problems and a detailed study based on CPT models would be very beneficial. (See [72] for the notion of proportional fairness considered in the classical setting of network resource allocation with EUT players.) We leave this for future work.

The rich structure of CPT towards capturing probabilistic sensitivity of the players provides additional flexibility in allocating resources based on individual preferences and needs. For example, our framework allows players to indicate their varying preferences towards different probabilities of winning as opposed to simply their valuation of winning (with certainty). In terms of markets, this implies that players can now compete for marginal probabilities of gaining resources. To appreciate the importance of this, consider the times when due to surge pricing many individuals with limited budget are capped out from the market. Often, this dissuades the smaller players from participating in these markets leading to less diverse markets. A lottery based approach would provide a chance for the smaller participants to remain in the markets even under peak-price conditions. Surely, they will

have to pay a premium rate to access the partial chance of winning, but they can do so by staying within their budget limits. For example, suppose the budget of an individual is \$100 and the current price of the item is \$2000. The individual however is ready to pay \$50 for a 2% chance of winning the item. This puts him in a competitive position as compared to an EUT player who values the item at \$2000 and has unlimited budget. (Because the EUT player would be ready to pay only \$40 for a 2% chance of winning the item.)

One method used to assign bands to users is auctions, since FCC found this method to be the most profitable as they earned millions of dollars through auctions. . . . A potential problem of this method is that smaller companies may be priced out of the market and unable to compete with large firms. This would reduce the number of points of view in the communications industry, which would violate one of the principles of the FCC, to protect the public interest. Another method used to allocate bands of frequencies was lotteries. Lotteries were used by the FCC in the 1980s. A benefit of lotteries was that it gave all parties a chance at winning, unlike auctions which favor parties with more money. By giving all parties a chance it was believed that it served the public interest better.

Network Security, Augsburg, Germany.

Our framework organically combines lotteries and auctions and it can help address the above features in a unified manner. A related problem associated with the use of lotteries is the verification that the announced lotteries are indeed implemented. Methods from cryptography can be used to achieve this. Algorithms developed in computer science will be of particular use. Alternatively, in a repeated interaction setting one could use the notion of calibration (as considered in Chapter 5) to check the consistency of the allocations with the proposed lotteries. For example, a customer who participates in such lotteries repeatedly would expect that among all the times she was promised a 10% lottery, she won the lottery with a frequency not too lower than 10%. This raises several interesting questions such as can the system operator deviate from the promised lotteries and still end up being calibrated. If yes, then it can potentially use such tactics to increase its revenue. There are other notions of calibration similar to the one considered in Chapter 5 and different metrics of calibratedness. One could ask similar questions with these notions providing another interesting direction for future work.

Finally, we consider some of the ethical considerations that relate to the use of lotteries which are often associated with negative connotations. For example, they are often perceived as preying on the emotional and psychological aspects of the players. As identified by Kahneman and Tversky, framing effects and emotional response by the individuals are a norm rather than an exception. Since there is no way of avoiding them, it is better we understand them and incorporate them in our models. Besides, offering lotteries and making use of richer frameworks as proposed in this thesis would allow for more flexibility in catering to the preferences of the individuals. Any statement regarding the social benefits of these general

methods would depend on further deliberation and on that particular instance. Nonetheless, it is definitely worth considering these lottery-based methods for resource allocation.

In this regard, recall the hyper-parameter k used in the discretization trick in Chapter 2. This way we restricted the probabilities offered to the players to be an integral multiple of $1/k$. For example, when $k = 100$, this amounts to showing integral percentages to the players. This is in contrast to the lotteries where the chance of winning is less than one in a million or lotteries where the true probabilities are obscured. Such lotteries often take advantage of the difficulty the players face in comprehending the true probabilities. By fixing k within reasonable limits depending on the situation, one could potentially avoid taking advantage of the players' limited rationality.

Another feature of market-based mechanisms that are designed to align with each individual's preferences is that it creates differentiated pricing and opportunity discrimination. This arises purely from the demand for resources and congestion in the market as well as the behavioral preferences and needs of the individuals. Any influence of factors such as race, color, religion, or sex will be accompanied by the approximations introduced by the use of AI to reduce complexities in our framework. This issue related to the use of AI is often referred to as bias in AI.

7.4 Conclusion

Combined efforts from fields such as game theory, behavioral economics, psychology, network economics, and computer science are needed to develop social systems for optimal allocation of resources, improved social welfare, and algorithmically assisted decision-making. In this thesis, we saw that CPT is an excellent tool for the study of *behavioral network economics*.

Bibliography

- [1] Mohammed Abdellaoui, Han Bleichrodt, and Hilda Kammoun. “Do financial professionals behave according to prospect theory? An experimental study”. In: *Theory and Decision* 74.3 (2013), pp. 411–429.
- [2] Marina Agranov and Pietro Ortoleva. “Stochastic choice and preferences for randomization”. In: *Journal of Political Economy* 125.1 (2017), pp. 40–68.
- [3] Maurice Allais. “Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école américaine”. In: *Econometrica: Journal of the Econometric Society* (1953), pp. 503–546.
- [4] Eitan Altman and Laura Wynter. “Equilibrium, games, and pricing in transportation and telecommunication networks”. In: *Networks and Spatial Economics* 4.1 (2004), pp. 7–21.
- [5] Robert J Aumann. “Correlated equilibrium as an expression of Bayesian rationality”. In: *Econometrica: Journal of the Econometric Society* (1987), pp. 1–18.
- [6] Robert J Aumann. “Subjectivity and correlation in randomized strategies”. In: *Journal of mathematical Economics* 1.1 (1974), pp. 67–96.
- [7] Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information*. MIT press, 1995.
- [8] Robert Aumann and Adam Brandenburger. “Epistemic conditions for Nash equilibrium”. In: *Econometrica: Journal of the Econometric Society* (1995), pp. 1161–1180.
- [9] Eduardo M Azevedo and Daniel Gottlieb. “Risk-neutral firms can extract unbounded profits from consumers with prospect theory preferences”. In: *Journal of Economic Theory* 147.3 (2012), pp. 1291–1299.
- [10] Nicholas C Barberis. “Thirty years of prospect theory in economics: A review and assessment”. In: *Journal of Economic Perspectives* 27.1 (2013), pp. 173–96.
- [11] Nicholas Barberis and Ming Huang. “Preferences with frames: a new utility specification that allows for the framing of risks”. In: *Journal of Economic Dynamics and Control* 33.8 (2009), pp. 1555–1576.
- [12] Yoram Barzel. “A theory of rationing by waiting”. In: *The Journal of Law and Economics* 17.1 (1974), pp. 73–95.

- [13] Manel Baucells, Martin Weber, and Frank Welfens. “Reference-point formation and updating”. In: *Management Science* 57.3 (2011), pp. 506–519.
- [14] Shlomo Benartzi and Richard H Thaler. “Myopic loss aversion and the equity premium puzzle”. In: *The quarterly journal of Economics* 110.1 (1995), pp. 73–92.
- [15] Dirk Bergemann and Stephen Morris. “Robust mechanism design”. In: *Econometrica* (2005), pp. 1771–1813.
- [16] David Blackwell. “An analog of the minimax theorem for vector payoffs”. In: *Pacific Journal of Mathematics* 6.1 (1956), pp. 1–8.
- [17] Robert F Bordley. “An intransitive expectations-based Bayesian variant of prospect theory”. In: *Journal of Risk and Uncertainty* 5.2 (1992), pp. 127–144.
- [18] Robert Bordley and Gordon B Hazen. “SSB and weighted linear utility as expected utility with suspicion”. In: *Management Science* 37.4 (1991), pp. 396–408.
- [19] Tilman Börgers and Daniel Krahmer. *An introduction to the theory of mechanism design*. Oxford University Press, USA, 2015.
- [20] Tilman Börgers and Taejun Oh. “Common Prior Type Spaces in Which Payoff Types and Belief Types are Independent”. In: *polar* 2011.2011b (2011).
- [21] John R Boyce. “Allocation of goods by lottery”. In: *Economic Inquiry* 32.3 (1994), pp. 457–476.
- [22] George W Brown. “Iterative solution of games by fictitious play”. In: *Activity analysis of production and allocation* 13.1 (1951), pp. 374–376.
- [23] Antonio Calvo-Armengol et al. “The set of correlated equilibria of 2×2 games”. In: *Work* (2006).
- [24] Colin F Camerer. “Prospect theory in the wild: Evidence from the field”. In: *Choices, Values, and Frames*. Contemporary Psychology. No. 47. American Psychology Association, Washington, DC, 2001, pp. 288–300.
- [25] Colin F Camerer and Teck-Hua Ho. “Violations of the betweenness axiom and non-linearity in probability”. In: *Journal of risk and uncertainty* 8.2 (1994), pp. 167–196.
- [26] Sabrina Gomez Canovas, Pierre Hansen, and Brigitte Jaumard. “Nash Equilibria from the correlated equilibria viewpoint”. In: *International Game Theory Review* 1.01 (1999), pp. 33–44.
- [27] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- [28] Deeparnab Chakrabarty, Nikhil Devanur, and Vijay V Vazirani. “New results on rationality and strongly polynomial time solvability in Eisenberg-Gale markets”. In: *International Workshop on Internet and Network Economics*. Springer. 2006, pp. 239–250.

- [29] J. T. Chang and D. Pollard. “Conditioning as Disintegration”. In: *Statistica Neerlandica* 51.3 (1997), pp. 287–317.
- [30] Alain Chateauneuf and Peter Wakker. “An axiomatization of cumulative prospect theory for decision under risk”. In: *Journal of Risk and Uncertainty* 18.2 (1999), pp. 137–145.
- [31] Yeon-Koo Che and Ian Gale. “Optimal design of research contests”. In: *American Economic Review* 93.3 (2003), pp. 646–671.
- [32] S Chew and Kenneth MacCrimmon. “Alpha-nu choice theory: an axiomatization of expected utility”. In: *University of British Columbia Faculty of Commerce working paper* 669 (1979).
- [33] Soo Hong Chew. “A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox”. In: *Econometrica: Journal of the Econometric Society* (1983), pp. 1065–1092.
- [34] Soo Hong Chew. “Axiomatic utility theories with the betweenness property”. In: *Annals of operations Research* 19.1 (1989), pp. 273–298.
- [35] Soo-hong Chew. *Implicit Weighted and Semi-Weighted Utility Theories, M-Estimators, and Non-Demand Revelation of Second-Price Auctions for an Uncertain Auctioned Object*. Johns Hopkins Univ., Department of Political Economy, 1985.
- [36] Vincent Conitzer and Tuomas Sandholm. “Computational criticisms of the revelation principle”. In: *Proceedings of the 5th ACM conference on Electronic commerce*. 2004, pp. 262–263.
- [37] Vincent P Crawford. “Equilibrium without independence”. In: *Journal of Economic Theory* 50.1 (1990), pp. 127–154.
- [38] M Cripps. *Extreme correlated and Nash equilibria in two-person games*. 1995.
- [39] Enrico G De Giorgi and Shane Legg. “Dynamic portfolio choice and asset pricing with narrow framing and probability weighting”. In: *Journal of Economic Dynamics and Control* 36.7 (2012), pp. 951–972.
- [40] Eddie Dekel. “An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom”. In: *Journal of Economic theory* 40.2 (1986), pp. 304–318.
- [41] Nadja Dwenger, Dorothea Kübler, and Georg Weizsäcker. “Flipping a coin: Theory and evidence”. In: (2012).
- [42] Sebastian Ebert and Philipp Strack. “Until the bitter end: on prospect theory in a dynamic context”. In: *American Economic Review* 105.4 (2015), pp. 1618–33.
- [43] Torstein Eckhoff. “Lotteries in allocative situations”. In: *Information (International Social Science Council)* 28.1 (1989), pp. 5–22.

- [44] Edmund Eisenberg and David Gale. “Consensus of subjective probabilities: The parimutuel method”. In: *The Annals of Mathematical Statistics* 30.1 (1959), pp. 165–168.
- [45] Fe S Evangelista and T ES Raghavan. “A note on correlated equilibrium”. In: *International Journal of Game Theory* 25.1 (1996), pp. 35–41.
- [46] Matthias Falkner, Michael Devetsikiotis, and Ioannis Lambadaris. “An overview of pricing concepts for broadband IP networks”. In: *IEEE Communications Surveys & Tutorials* 3.2 (2000), pp. 2–13.
- [47] Peter C Fishburn. *Nonlinear preference and utility theory*. 5. Johns Hopkins University Press Baltimore, 1988.
- [48] Peter C Fishburn and Gary A Kochenberger. “Two-piece von Neumann-Morgenstern utility functions”. In: *Decision Sciences* 10.4 (1979), pp. 503–518.
- [49] Dean P Foster and Rakesh V Vohra. “Asymptotic calibration”. In: *Biometrika* 85.2 (1998), pp. 379–390.
- [50] Dean P Foster and Rakesh V Vohra. “Calibrated learning and correlated equilibrium”. In: *Games and Economic Behavior* 21.1-2 (1997), pp. 40–55.
- [51] Dean P Foster and H Peyton Young. “Regret testing: Learning to play Nash equilibrium without knowing you have an opponent”. In: *Theoretical Economics* 1.3 (2006), pp. 341–367.
- [52] Drew Fudenberg and David Levine. “Learning in games”. In: *European economic review* 42.3-5 (1998), pp. 631–639.
- [53] Drew Fudenberg and David K Levine. “Consistency and cautious fictitious play”. In: *Journal of Economic Dynamics and Control* 19.5-7 (1995), pp. 1065–1089.
- [54] Sanjeev Goyal. *Connections*. Princeton University Press, 2012.
- [55] Faruk Gul. “A theory of disappointment aversion”. In: *Econometrica: Journal of the Econometric Society* (1991), pp. 667–686.
- [56] James Hannan. “Approximation to Bayes risk in repeated play”. In: *Contributions to the Theory of Games* 3 (1957), pp. 97–139.
- [57] Bruce GS Hardie, Eric J Johnson, and Peter S Fader. “Modeling loss aversion and reference dependence effects on brand choice”. In: *Marketing science* 12.4 (1993), pp. 378–394.
- [58] John C Harsanyi. “Games with incomplete information played by “Bayesian” players, I–III Part I. The basic model”. In: *Management science* 14.3 (1967), pp. 159–182.
- [59] Sergiu Hart. “Adaptive heuristics”. In: *Econometrica* 73.5 (2005), pp. 1401–1430.
- [60] Sergiu Hart and Andreu Mas-Colell. “A simple adaptive procedure leading to correlated equilibrium”. In: *Econometrica* 68.5 (2000), pp. 1127–1150.

- [61] Aanund Hylland and Richard Zeckhauser. “The efficient allocation of individuals to positions”. In: *Journal of Political economy* 87.2 (1979), pp. 293–314.
- [62] Matthew O Jackson. *Social and economic networks*. Princeton university press, 2010.
- [63] Matthew O Jackson. *The human network: how your social position determines your power, beliefs, and behaviors*. Vintage, 2019.
- [64] Kamal Jain and Vijay V Vazirani. “Eisenberg–Gale markets: Algorithms and game-theoretic properties”. In: *Games and Economic Behavior* 70.1 (2010), pp. 84–106.
- [65] Steven J Kachelmeier and Mohamed Shehata. “Examining risk preferences under high monetary incentives: Experimental evidence from the People’s Republic of China”. In: *The American economic review* (1992), pp. 1120–1141.
- [66] Daniel Kahneman. *Thinking, fast and slow*. Macmillan, 2011.
- [67] Daniel Kahneman and Amos Tversky. “Prospect theory: An analysis of decision under risk”. In: *Econometrica: Journal of the Econometric Society* 47.2 (1979), pp. 263–291.
- [68] Daniel Kahneman and Amos Tversky. “Prospect theory: An analysis of decision under risk”. In: *Handbook of the fundamentals of financial decision making: Part I*. World Scientific, 2013, pp. 99–127.
- [69] Shizuo Kakutani. “A generalization of Brouwer’s fixed point theorem”. In: *Duke mathematical journal* 8.3 (1941), pp. 457–459.
- [70] Dileep Kalathil, Vivek S Borkar, and Rahul Jain. “Approachability in Stackelberg stochastic games with vector costs”. In: *Dynamic games and applications* 7.3 (2017), pp. 422–442.
- [71] Edi Karni and Zvi Safra. “Dynamic consistency, revelations in auctions and the structure of preferences”. In: *The Review of Economic Studies* 56.3 (1989), pp. 421–433.
- [72] Frank Kelly. “Charging and rate control for elastic traffic”. In: *European Transactions on Telecommunications* 8.1 (1997), pp. 33–37.
- [73] Frank P Kelly, Aman K Maulloo, and David KH Tan. “Rate control for communication networks: shadow prices, proportional fairness and stability”. In: *Journal of the Operational Research Society* 49.3 (1998), pp. 237–252.
- [74] Kerim Keskin. “Equilibrium Notions for Agents with Cumulative Prospect Theory Preferences”. In: *Decision Analysis* 13.3 (2016), pp. 192–208.
- [75] David M Kreps. *Game theory and economic modelling*. Oxford University Press, 1990.
- [76] Richard J La and Venkat Anantharam. “Utility-based rate control in the Internet for elastic traffic”. In: *IEEE/ACM Transactions on Networking (TON)* 10.2 (2002), pp. 272–286.
- [77] James M Lattin and Randolph E Bucklin. “Reference effects of price and promotion on brand choice behavior”. In: *Journal of Marketing research* 26.3 (1989), pp. 299–310.

- [78] Philip Leclerc. “Prospect theory preferences in noncooperative game theory”. PhD thesis. Virginia Commonwealth University, 2014.
- [79] Xiaojun Lin, Ness B Shroff, and Rayadurgam Srikant. “A tutorial on cross-layer optimization in wireless networks”. In: *IEEE Journal on Selected areas in Communications* 24.8 (2006), pp. 1452–1463.
- [80] Qingmin Liu. “On redundant types and Bayesian formulation of incomplete information”. In: *Journal of Economic Theory* 144.5 (2009), pp. 2115–2145.
- [81] James G March. “Variable risk preferences and adaptive aspirations”. In: *Journal of Economic Behavior & Organization* 9.1 (1988), pp. 5–24.
- [82] James G March and Zur Shapira. “Variable risk preferences and the focus of attention.” In: *Psychological review* 99.1 (1992), p. 172.
- [83] Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. *Microeconomic theory*. Vol. 1. Oxford university press New York, 1995.
- [84] Jeonghoon Mo and Jean Walrand. “Fair end-to-end window-based congestion control”. In: *IEEE/ACM Transactions on Networking* 8.5 (2000), pp. 556–567.
- [85] Benny Moldovanu and Aner Sela. “The optimal allocation of prizes in contests”. In: *American Economic Review* 91.3 (2001), pp. 542–558.
- [86] Dilip Mookherjee and Masatoshi Tsumagari. “Mechanism design with communication constraints”. In: *Journal of Political Economy* 122.5 (2014), pp. 1094–1129.
- [87] John Morgan. “Financing public goods by means of lotteries”. In: *The Review of Economic Studies* 67.4 (2000), pp. 761–784.
- [88] Roger B Myerson. “Comments on “Games with Incomplete Information Played by ‘Bayesian’ Players, I–III Harsanyi’s Games with Incomplete Information””. In: *Management Science* 50.12_supplement (2004), pp. 1818–1824.
- [89] Roger B Myerson. *Game theory*. Harvard university press, 2013.
- [90] Roger B Myerson. “Incentive compatibility and the bargaining problem”. In: *Econometrica: journal of the Econometric Society* (1979), pp. 61–73.
- [91] Roger B Myerson. “Multistage games with communication”. In: *Econometrica: Journal of the Econometric Society* 54.2 (1986), pp. 323–358.
- [92] Roger B Myerson. “Optimal auction design”. In: *Mathematics of operations research* 6.1 (1981), pp. 58–73.
- [93] Roger B Myerson. “Optimal coordination mechanisms in generalized principal–agent problems”. In: *Journal of mathematical economics* 10.1 (1982), pp. 67–81.
- [94] Anna Nagurney. *Network economics: A variational inequality approach*. Vol. 10. Springer Science & Business Media, 2013.
- [95] Anna Nagurney. *Supply chain network economics: dynamics of prices, flows and profits*. Edward Elgar Publishing, 2006.

- [96] Anna Nagurney et al. “Sustainable transportation networks”. In: *Books* (2000).
- [97] Anna Nagurney and Qiang Qiang. *Fragile networks: identifying vulnerabilities and synergies in an uncertain world*. John Wiley & Sons, 2009.
- [98] Anna Nagurney and Stavros Siokos. *Financial networks: Statics and dynamics*. Springer Science & Business Media, 2012.
- [99] John Nash. “Non-cooperative games”. In: *Annals of Mathematics* (1951), pp. 286–295.
- [100] John F Nash. “Equilibrium points in n-person games”. In: *Proceedings of the national academy of sciences* 36.1 (1950), pp. 48–49.
- [101] Robert Nau, Sabrina Gomez Canovas, and Pierre Hansen. “On the geometry of Nash equilibria and correlated equilibria”. In: *International Journal of Game Theory* 32.4 (2004), pp. 443–453.
- [102] Michael Nwogugu. “A further critique of cumulative prospect theory and related approaches”. In: *Applied mathematics and computation* 179.2 (2006), pp. 451–465.
- [103] Paul A Pavlou and Angelika Dimoka. “The nature and role of feedback text comments in online marketplaces: Implications for trust building, price premiums, and seller differentiation”. In: *Information Systems Research* 17.4 (2006), pp. 392–414.
- [104] Vianney Perchet. “Calibration and internal no-regret with random signals”. In: *International Conference on Algorithmic Learning Theory*. Springer. 2009, pp. 68–82.
- [105] Soham R Phade and Venkat Anantharam. “Black-Box Strategies and Equilibrium for Games with Cumulative Prospect Theoretic Players”. In: *arXiv preprint arXiv:2004.09592* (2020).
- [106] Soham R Phade and Venkat Anantharam. “Learning in Games with Cumulative Prospect Theoretic Preferences”. In: *Dynamic Games and Applications* (2021). DOI: 10.1007/s13235-021-00398-9. URL: <https://doi.org/10.1007/s13235-021-00398-9>.
- [107] Soham R Phade and Venkat Anantharam. “Mechanism Design for Cumulative Prospect Theoretic Agents: A General Framework and the Revelation Principle”. In: *arXiv preprint arXiv:2101.08722* (2021).
- [108] Soham R Phade and Venkat Anantharam. “On the Geometry of Nash and Correlated Equilibria with Cumulative Prospect Theoretic Preferences”. In: *Decision Analysis* 16.2 (2019), pp. 142–156.
- [109] Soham R Phade and Venkat Anantharam. “Optimal Resource Allocation over Networks via Lottery-Based Mechanisms”. In: *arXiv preprint* (2018).
- [110] Thierry Post et al. “Deal or no deal? decision making under risk in a large-payoff game show”. In: *American Economic Review* 98.1 (2008), pp. 38–71.

- [111] Balaji Prabhakar. “Designing large-scale nudge engines”. In: *ACM SIGMETRICS Performance Evaluation Review*. Vol. 41. 1. ACM. 2013, pp. 1–2.
- [112] Drazen Prelec. “A “Pseudo-endowment” effect, and its implications for some recent nonexpected utility models”. In: *Journal of Risk and Uncertainty* 3.3 (1990), pp. 247–259.
- [113] Drazen Prelec. “The probability weighting function”. In: *Econometrica* (1998), pp. 497–527.
- [114] John Quiggin. “A theory of anticipated utility”. In: *Journal of economic behavior & organization* 3.4 (1982), pp. 323–343.
- [115] John Quiggin. “On the optimal design of lotteries”. In: *Economica* (1991), pp. 1–16.
- [116] “Review of the Internet traffic management practices of Internet service providers”. In: *Telecom. Reg. Policy CRTC 2009-657* (2009).
- [117] Ralph Tyrell Rockafellar. *Convex analysis*. Princeton university press, 2015.
- [118] Alvin E Roth and Elliott Peranson. “The redesign of the matching market for American physicians: Some engineering aspects of economic design”. In: *American economic review* 89.4 (1999), pp. 748–780.
- [119] David Schmeidler. “Subjective probability and expected utility without additivity”. In: *Econometrica: Journal of the Econometric Society* (1989), pp. 571–587.
- [120] Paul JH Schoemaker. “The expected utility model: Its variants, purposes, evidence and limitations”. In: *Journal of economic literature* (1982), pp. 529–563.
- [121] Jonathan Shalev. “Loss aversion equilibrium”. In: *International Journal of Game Theory* 29.2 (2000), pp. 269–287.
- [122] Carl Shapiro and Hal R. Varian. *Information Rules : A Strategic Guide to the Network Economy*. Harvard Business Review Press, 1998. URL: <https://search-ebshost-com.libproxy.berkeley.edu/login.aspx?direct=true&db=nlebk&AN=35060&site=ehost-live>.
- [123] Barry Sopher and J Mattison Narramore. “Stochastic choice and consistency in decision making under risk: An experimental study”. In: *Theory and Decision* 48.4 (2000), pp. 323–350.
- [124] Chris Starmer. “Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk”. In: *Journal of economic literature* 38.2 (2000), pp. 332–382.
- [125] Peter Stone. “Why lotteries are just”. In: *Journal of Political Philosophy* 15.3 (2007), pp. 276–295.
- [126] Grant A Taylor, Kevin KK Tsui, and Lijing Zhu. “Lottery or waiting-line auction?”. In: *Journal of Public Economics* 87.5-6 (2003), pp. 1313–1334.

- [127] Amos Tversky and Daniel Kahneman. “Advances in prospect theory: Cumulative representation of uncertainty”. In: *Journal of Risk and uncertainty* 5.4 (1992), pp. 297–323.
- [128] Yannick Viossat. *The geometry of Nash equilibria and correlated equilibria and a generalization of zero-sum games*. Tech. rep. SSE/EFI Working Paper Series in Economics and Finance, 2006.
- [129] Rakesh V Vohra. *Mechanism design: A linear programming approach*. Vol. 47. Cambridge University Press, 2011.
- [130] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2007.
- [131] Peter Wakker. “Separating marginal utility and probabilistic risk aversion”. In: *Theory and decision* 36.1 (1994), pp. 1–44.
- [132] Peter P Wakker. *Prospect theory: For risk and ambiguity*. Cambridge university press, 2010.
- [133] Jiantao Wang et al. “Cross-layer optimization in TCP/IP networks”. In: *IEEE/ACM Transactions on Networking (TON)* 13.3 (2005), pp. 582–595.
- [134] Steven R Williams. *Communication in mechanism design: A differential approach*. Cambridge University Press, 2008.
- [135] Menahem E Yaari. “The dual theory of choice under risk”. In: *Econometrica: Journal of the Econometric Society* (1987), pp. 95–115.
- [136] H Peyton Young. *Strategic learning and its limits*. OUP Oxford, 2004.

Index

- 2x2 game, 57, 103
- action, 50, 96
 - mixture, 78
 - randomization, 96
- action profile, 50
- agent
 - see also* player 136
- agent-network decomposition, 14
- Allais paradox, 7
- allocation (mechanism design), 141
- allocation choice function, 142
- allocation function, 137, 144
 - mediated, 137, 157
- allocation-outcome mapping, 139, 141, 147
- Artificial Intelligence (AI), 192
 - Machine Learning (ML), 193
 - Reinforcement Learning (RL), 193
- assessment, 96
- auction, 20
 - second-price sealed bid, 16
- Bayes-Nash equilibrium, 137
 - mechanism design, 145, 158
- Bayesian game, 137
- belief, 77
 - conjecture, 77
- belief-dominant strategy, 146
- best reaction, 96
 - map, 96
- best response
 - action set, 77
- best response black-box strategy set, 78
- best response strategy set
 - for mechanism design, 145, 158
- betweenness, 15, 71, 74
 - weak, 74
- big data analytics, 192
- black-box strategies, 15
- black-box strategy, 72
- black-box strategy Nash equilibrium, 73
- Blackwell approachability, 95
- Braess' paradox, 2
- budget signals, 14, 28
- calibrated, 97
- calibrated learning, 95
- capacity constraints, 21
- correlated equilibrium, 50
- CPT correlated equilibrium, 51
- CPT features
 - probability weighting functions, 9
 - reference point, 9
 - value function, 9
- CPT Nash equilibrium, 52
- cumulative prospect theory (CPT), 7, 9
 - CPT preferences, 11
 - CPT value, 9
 - decision weights, 10
- diminishing sensitivity to returns, 9
- direct allocation function, 148
- direct mechanism, 148
- direct mediated mechanism, 158
- discretization trick, 24, 171, 196
- dominant strategy

- mechanism design, 145
- dominated strategy, 59
- endowment effect, 3
- ex-ante utility, 22
- expected utility theory (EUT), 7
 - EUT preferences, 7
 - expected utility, 7
 - utility function, 7
- games with communication, 95
- implementability, 140
 - Bayes-Nash equilibrium, 145
 - belief-dominant equilibrium, 146
 - dominant equilibrium, 145
 - truthful, 148
- independence axiom, 71, 74
- learning in games, 94
 - calibrated learning, 95
 - fictitious play, 111
 - follow the perturbed leader, 111
- lottery
 - pointwise dominance, 53
 - similarly ranked, 53
- lottery pricing, 14, 27
- lottery, prospect, 7
- lottery-based allocation
 - optimal structure, 36
- lottery-based allocations, 20
- mechanism design, 17
- mediated CPT correlated equilibrium, 16, 102
- mediated CPT Nash equilibrium, 101
- mediated game, 95, 101
- mediated mechanism, 137
 - direct, 158
- mediator distribution, 101
 - mechanism design, 137, 157
- message set, 137, 157
- mixed black-box strategy Nash equilibrium, 73
- mixed Nash equilibrium, 73
- Nash equilibrium, 50, 72
- network externality, 3
- network problem, 23, 27
- normal form game, 50
- NP-hard, 33
- nudge, 17
- outcome, 7
 - for mechanism design, 141
- outcome choice function, 142
- payoff function, 50
- permutation, 9, 21
 - optimum, 30
 - structure, 25
- player, 1, 50, 141
- probabilistic sensitivity, 9, 21
- prospect theory (PT), 13
- publicly mediated mechanism, 160
- pure Nash equilibrium, 73
- rank dependent utility (RDU), 13, 21, 76
- regret, 53, 94
 - CPT regret, 106
- repeated game
 - action play, 96
 - assessment, 96
 - empirical distribution of action play, 95
 - history, 96
 - stage game, 96
- revelation principle, 16, 95, 136, 137, 148, 160
- signal set, 137, 144
- social choice function, 140, 142
- strategy
 - mechanism design, 144
 - normal form game, 50
 - randomized, 101
 - repeated game, 96
- strict monotonicity, 13, 54, 56

- strict stochastic dominance, 12, 53, 151
- system operator, 136
- system problem, 23
 - average *see also* relaxation 33
 - duality gap, 33
 - equilibrium, 28
 - relaxation, 31
- traffic shaping, 4
- Transmission Control Protocol (TCP), 4, 14
- truthful implementation, 148
- type, 136
- user problem, 23, 27
- user-network decomposition
 - see also* agent-network decomposition
 - 23
- utilitarian, 22
- utility function
 - for mechanism design, 143