

Reasoning about (strategic) decisions under uncertainty

Citation for published version (APA):

Nauerz, C. T. (2016). *Reasoning about (strategic) decisions under uncertainty*. [Doctoral Thesis, Maastricht University]. Universitaire Pers Maastricht. <https://doi.org/10.26481/dis.20160921cn>

Document status and date:

Published: 01/01/2016

DOI:

[10.26481/dis.20160921cn](https://doi.org/10.26481/dis.20160921cn)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Reasoning About (Strategic) Decisions Under Uncertainty

Christian Nauerz

© Christian Nauerz, Maastricht 2016.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form, or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission in writing from the author.

This book was typeset by the author using L^AT_EX.

Published by Universitaire Pers Maastricht.

ISBN: 978-94-6159-597-3

Printed in The Netherlands by Datawyse Maastricht

Reasoning About (Strategic) Decisions Under Uncertainty

DISSERTATION

to obtain the degree of Doctor
at Maastricht University,
on the authority of the Rector Magnificus
Prof. dr. Rianne M. Letschert,
in accordance with the decision of the Board of Deans,
to be defended in public
on Wednesday, September 21, 2016, at 16.00 hours

by

Christian Tobias Nauerz



Supervisor

Prof. Dr. Hans Peters

Co-supervisor

Dr. Andrés Perea

Dr. Elias Tsakas

Assessment Committee

Prof. Dr. Arno Riedl (Chair)

Prof. Dr. Peter Wakker (Erasmus University, Rotterdam)

Dr. Christian Bach (University of Liverpool, Liverpool)

Dr. Matthew Embrey (University of Sussex, Brighton)

This research was financially supported by the Graduate School of Business and Economics (GSBE).

To my parents, Jutta & Jürgen.

Acknowledgments

The journey towards a Ph.D. degree is a turbulent experience. In the short-term, the connection between the input in terms of hours spend on a given topic and the output of research results can be very noisy. At times, trivial intuitions are difficult to translate into a rigorous language, whereas more involved results follow effortlessly. Sometimes success depends simply on taking the right perspective towards a problem. Thus, the challenge is often not to have an ingenious insight but to keep the overview of the problem and to be resilient to failure. This can lead to a rapid fluctuation between euphuism and devastation. In these situations, I could always count on the people around me. They gave me the support, to keep on going in the face of failure and kept my feet grounded when things went well.

First of all, I would like to thank my parents for the unconditional and loving support throughout the years. From my childhood on, they provided an environment that allowed me to grow and thrive, with the occasional nudge to keep on going.

Thank you Andrés and Elias for being my supervisors. I am honored that you took all the time and effort to support me during this journey. You were always available for a short coffee brake to discuss some research ideas. Many times you saved me going down the wrong path and spending countless hours on a hopeless pursuit. You patently listened to all my ideas and gave me valuable feedback. I cannot image a better team of supervisors.

The quality of research often depends on the right perspective. Thus, I am indebted to all my colleagues who were always available for discussions. Especially, I would like to thank Christian Bach, Matthew Embrey, János Flesch, Ronald Peters, and Martin Strobel for the time they took to discuss research ideas.

I am also grateful to David Ahn for inviting me to Berkeley. It was a truly

great experience both professionally and socially. I had great discussions with Shachar Kariv, Yoram Halevy, and of course David. I am also thankful for meeting so many interesting and inspiring people inside and outside of the university. It was a time I will always cherish.

Also, I am thankful that I was surrounded by great friends. Frauke and Marion, it was a great experience to write a paper together with the two of you. The ups and downs of the process are more easily digested when jokes lighten the mood and you can build on each other's ideas. Anna and Angie, I am glad that we could talk to each other and share our successes and failures. Our dinners that often ended in philosophical discussions about science is something I will always remember. I was lucky that I always had great office mates: Artem, Daria, and Verena. Exchanging our experience and sharing some laughs made often made my day. A thank you also goes to my colleagues at the KE department and the ROA girls. It was a lot of fun sharing birthday cakes and drinks with you guys.

A special thank you goes to my friends from my study time at Ingolstadt, Sophie, Marco and Walter. Without you, I wouldn't be who I am. Thank you for always being there.

Last but not least, I would like to thank the reading committee for their time and energy they spend assessing my thesis.

Contents

Acknowledgements	viii
1 Introduction	1
2 Beyond Coincidence: The Reasoning Process Underlying Utility Proportional Beliefs	9
2.1 Epistemic Characterization	11
2.2 Computing Beliefs Directly	15
2.3 Cognitive Characterization	18
2.4 Example	22
2.5 Experimental Evidence	24
2.6 Conclusion	28
2.A Appendix - Proofs	29
3 Identifying Strategic Thinking in One-Shot Games	31
3.1 Concepts	34
3.2 Experimental Method	37
3.2.1 Procedures	37
3.2.2 Games	38
3.2.3 Belief Elicitation and Incentives	40
3.3 Results	43
3.3.1 Descriptive Evidence	43
3.3.2 Performance Against Random Choice	46
3.3.3 Performance Against Perfect Consistency	47
3.3.4 Response Times	48
3.3.5 Best Responses	50
3.4 Discussion	52

3.5	Concluding Remarks	55
3.A	Appendix - Logit-QRE as Differences in Utilities	56
3.B	Appendix - Description of Games	56
3.C	Appendix - Games	59
3.D	Appendix - Robustness Classification	62
3.E	Appendix - Instructions and Control Questions	62
4	Local Prior Expected Utility: A Basis for Utility Representations Under Uncertainty	71
4.1	A Utility Function for Acts	73
4.2	Gradient of ψ	74
4.3	Representation for Local Prior Expected Utility	76
4.4	Representation for Classical Models of Decision Making Under Uncertainty	79
4.4.1	Maximin Expected Utility	79
4.4.2	Choquet Expected Utility	83
4.4.3	Anscombe-Aumann	88
4.5	Discussion	89
4.5.1	Characterizations of the Set of Priors	89
4.5.2	Further Weakening of the C-Independence Axiom	90
2.A	Appendix - Proofs Section 4.1	93
2.B	Appendix - Proofs Section 4.2	95
2.C	Appendix - Proofs Section 4.3	96
2.D	Appendix - Proofs Section 4.4	101
2.E	Appendix - Proofs Section 3.4	106
	Bibliography	109
	Valorization	115
	Nederlandse samenvatting	121
	Curriculum Vitae	127

Chapter 1

Introduction

Research in economics is about understanding and evaluating decisions of economic agents, a question that needs to be investigated theoretically and empirically. I believe that the two approaches are strongly interconnected and progress is only possible with advances in both areas. To accomplish this task, we first have to define the optimal behavior in a given situation. This is called the normative approach. The descriptive approach, on the other hand, focuses on understanding peoples' actual behavior. This thesis takes a descriptive perspective in both the theoretical and empirical investigations since the normative aspect of the problems at hand are well understood. However, at times normative aspects will enter the picture as we will have to reconsider our definition of rational behavior in some situations. A great example is the Ellsberg Paradox, which cannot be explained by standard models but the behavior it evokes can also not be considered to be irrational. Therefore, this thesis attempts to provide descriptive insights but will on the way also touch some normative questions.

Decision theory, a subfield within economics, investigates assumptions on the agents' rationality in order to explain their economic behavior. These can be static or dynamic decision problems including risk and uncertainty. Strategic uncertainty is, however, excluded from the analysis and treated within another stream of the literature called game theory. Even though, both fields are in principle tightly connected, there was no common foundation until the seminal work of Harsanyi (1967-68), Armbruster and Böge (1979), Brandenburger and Dekel (1987), Aumann (1987), and Tan and Werlang (1988). A major problem when dealing with strategic uncertainty is the complexity that arises. Suppose you

have to make a choice where the outcome is also dependent on another person's choice. Of course, you are uncertain about the other person's choice. Yet, this does not exhaust the uncertainty you are facing. To verify that your belief about the other person's choice is justified, you also need to consider his belief about your choice, giving rise to a new level of uncertainty. And again, to verify your belief about the other person's beliefs you need to consider what you believe the other person believes you believe about him. Continuing in this fashion adds infinitely many layers of uncertainty, an infinite belief hierarchy. Fortunately, Harsanyi (1967-68), Armbruster and Böge (1979), Böge and Eisele (1979), Mertens and Zamir (1985), and Brandenburger and Dekel (1987) showed that infinite belief hierarchies can be encoded in a single mathematical object called type. This theoretical achievement allows us to investigate non-strategic and strategic decision problems with the same class of models. The field concerned with the analysis of strategic decision making is called epistemic game theory. For a detailed overview of the field see Perea (2012).

After Nash (1950, 1951) introduced his famous equilibrium concept, game theory was mainly concerned with equilibrium analyses and refinements thereof. Epistemic game theory, on the other hand, does not rest on the equilibrium assumption. It rests on the idea of belief in the opponents rationality (however rationality might be defined). Let us consider the standard notion of rationality, where an agent is rational if he does not play a strictly dominated strategy. Now suppose you are playing a game against an opponent with a single dominated strategy. Then if you believe in the opponent's rationality, you cannot believe that he will be playing the dominated strategy. Thus, your belief in your opponent's rationality restricts the set of beliefs that you can hold. This short example illustrates game theoretic analyses within epistemic game theory. Of course, classical game theory and epistemic game theory are tightly connected. Much work in epistemic game theory is aimed at providing a deeper understanding of existing concepts in classical game theory such as the seminal work of Brandenburger and Dekel (1987), Aumann and Brandenburger (1995), and Tan and Werlang (1988).

A natural next step, when attempting to improve descriptive theories, is to ask about the reasoning process that generates the outcomes suggested by a concept. Yet, this question has not received a lot of attention in the realm of game theory. One reason being that this approach is outside of the toolbox of classical economics. Economists rely heavily on choice data since it is readily

available in most contexts, whereas even data such as beliefs is hard to obtain. This restriction, however, often leads to accepting models as valid if they match a small amount of choice data. An example from decision theory is discussed in Rubinstein (2003). Rubinstein criticizes hyperbolic discounting and shows that there are multiple and indistinguishable models explaining the data. By means of an experiment, he illustrates that there are situations of dynamic inconsistency that cannot be captured by the hyperbolic discounting model. The suggested reasoning process, however, has no problem of explaining these supposedly inconsistent choices. In general the description of a reasoning process can help to make an intuitive judgement about the likelihood that people might actually apply it. Take Nash equilibrium, for example. Nash equilibrium is the predominant concept in game theory. It describes the set of choices in a strategic situation where none of the participants has an incentive to deviate from his choice. This is clearly a useful concept when we consider strategic situations that are repeated often. Then it is likely that over the course of time participants' choices converge to a point where non of them has an incentive to deviate (if these points exist). Nevertheless, in a one-shot interaction it is unlikely that a non-game theorist tries to find a point where neither he nor his opponents have an incentive to change their choice. Two objections are immediate: First, what is the motivation to look for such a point? It does not necessarily maximize utility nor is such a choice necessarily unique. Second, the set of Nash equilibria is non-trivial to compute and even when we consider two participants where each has three choices, the calculations necessary to find the set of Nash equilibria can hardly be performed without pen and paper. However, there is a more severe conceptual hurdle that might not be immediate. As discussed in Perea (2012), in two player games Nash equilibrium requires players to express common belief in rationality as well as to believe that their opponent has correct beliefs and believe that their opponent believes that they themselves have correct beliefs.¹ It is unlikely, though, to have correct beliefs about an unknown opponent's actions in an unfamiliar strategic interaction without prior learning or communication. This example shows how insights into solution concepts can help us to understand their features and judge whether or not it is likely that people will apply these thought patterns when making choices in strategic settings. Thus I argue,

¹Note that these insights are based on Perea (2007) and can be extended to the n -player case by imposing slightly stronger assumptions. Similar conditions have been presented in Brandenburger and Dekel (1989), Aumann and Brandenburger (1995), Polak (1999), and Asheim (2006).

that when we attempt to understand peoples' strategic decision making, the underlying reasoning process should be intuitive and feasible.

In Chapter 2 I provide such a reasoning process of a game theoretic solution concept called utility proportional beliefs (Bach and Perea, 2014). The core idea of the concept is that the beliefs one assigns to his opponent's choices are proportional to the expected utility induced by the opponent's beliefs about one's own choices. Compared to other concepts, this idea ensures that beliefs are meaningful as they directly reflect the goodness of the opponent's choices. The assumption that agents hold utility proportional beliefs is, however, not sufficient. It would lead again to the problem that these utility proportional beliefs are not supported by higher order beliefs. To find the supporting higher order beliefs, we assume that the agent believes that his opponent also has utility proportional beliefs, that the agent believes that his opponent believes that he holds utility proportional beliefs, that the agent believes that his opponent believes that he believes that his opponent holds utility proportional beliefs, and so on. As before this leads to an infinite regress. A point where this infinite regress is satisfied for all agents is when there is common belief in utility proportional beliefs. Bach and Perea show that the beliefs satisfying common belief in utility proportional beliefs can be found by iterative elimination of all non-utility proportional beliefs from the set of possible beliefs.

When we consider the proposed elimination procedure, this concept does not seem to relate to an intuitive reasoning process. Finding non-utility proportional beliefs in the set of all possible beliefs is certainly not a natural approach for non-game theorists when looking at strategic decisions. Therefore, we take another approach. Bach and Perea find that the beliefs an agent can hold in strategic situations that do not involve more than two agents are unique under common belief in utility proportional beliefs. I used this insight to find a closed-form solution to obtain the beliefs that are possible under common belief in utility proportional beliefs. It turns out that this solution can be interpreted as a reasoning process leading to the beliefs agents can hold under common belief in utility proportional beliefs. The process resembles two intuitive properties. (1) Initial Belief: The agents' reasoning process starts with an initial belief about the opponent, which is the basis for the further belief formation process. It considers the goodness of the other agent's choices according to his average utility. (2) Interactive Reasoning Procedure: Within the reasoning process agents' change perspectives within every reasoning step. Hence, they consider the other agent's incentive

structure and their own alternately. Understanding this process allows us to formulate hypotheses about the time it takes to form certain beliefs. Using the data set generated by the experiment described in Chapter 3, I show that these hypotheses cannot be rejected for subjects who state utility proportional beliefs. Thus, these findings strengthen the trust in the concept of utility proportional beliefs beyond the coinciding beliefs. It shows that even predictions about the process of the belief formation can be verified.

Chapter 3 does not zoom in as deep as Chapter 2 but asks a similar question. Instead of investigating a single concept, Frauke Meyer, Marion Collewet, and I are interested in the heterogeneity between subjects with respect to their strategic thinking. The concepts we consider are quantal response equilibrium (QRE) (McKelvey and Palfrey, 1995), noisy level- k (Lk) (Stahl and Wilson, 1994, 1995), and utility proportional beliefs (UPB) (Bach and Perea, 2014). QRE is popular in the literature and has the convenient property that it is a noisy version of Nash equilibrium, the gold-standard of game theory. In particular, quantal response equilibria that are achieved in the limit, when the noise goes to zero, are Nash equilibria. Thus, we can test noisy equilibrium notions as well as Nash equilibrium by considering QRE. Models of level- k thinking are mainly motivated by Nagel's (1995) findings in the beauty contest game. The idea is that level-0 agents choose all choices with equal probability. Level-1 agents then best respond to level-0 agents' choices, i.e., choose such that they maximize their payoff given the choice of level-0 agents. Level-2 agents, again, best respond to the choice of level-1 agents, and so on. Nagel shows in her experiment that these types are robustly observed in beauty-contest games. Later contributions to the literature also found level- k agents in games other than the beauty contest. Lastly, we tested utility proportional beliefs as it has not been tested experimentally and has promising theoretical properties.

We elicit the subjects' beliefs about their opponent's choices. Beliefs might not be available in field data but can help us to gain more insights from lab data and therefore refine our intuition about the meaning of field data. The advantage of beliefs is that they provide deeper insights into subjects' strategic thinking. A single choice only tells us that an agent prefers the choice but not why and how strongly. Take the simple case where his belief induces almost equal expected utility of the choices at hand. Then his choice might not convey much information because the agent might take a shortcut and simply pick a choice at random after discovering that his potential loss from a mistake will be minimal. How-

ever, his belief informs us about his exact assessment of the situation. Since we believe that even if a subject consistently follows a single strategic approach it is unlikely that the noise influencing his beliefs will be constant across all games, we simply check if a subject's beliefs fall within the area of predicted beliefs for all possible noise levels for a given concept. Hence, we can simply check whether a belief falls within this area to conclude that a given choice is coherent with a model. However, this approach would favor the trivial theory that predicts the whole area of possible beliefs. Selten and Krischker's (Selten and Krischker, 1983) measure of predictive success (MPS) takes the predicted area into account when judging a concept's predictive power. This approach favors efficient concepts, i.e. concepts which predict well while only suggesting a small set of possible beliefs. Consequently, it selects the model that predicts a subject's beliefs most efficiently.

Using Selten and Krischker's MPS, we investigate two questions: 1) Does the concept with the highest measure of predictive success predict a subject's choice better than random choice? 2) Does a subject state beliefs as if he is thinking according to a given concept but makes some mistakes? We find that about 80 percent of the subjects' beliefs can be described to some degree by either level-k, QRE, or UPB, where level-k describes the strategic thinking of the majority of subjects. About 45 percent of the subjects state beliefs as if they reason according to a fixed concept and make some mistakes. Again, level-k describes the thinking of more subjects than any other concept. These results are remarkable as the majority of our games does not have a dominated strategy and therefore allows all possible beliefs under the standard rationality assumption. Therefore, our results present further evidence for the well-supported level-k models.

Finally, in Chapter 4 together with Andrés Perea I investigate economic decisions in cases where the decision maker has (possibly) more than one belief about the world. This view is motivated by the famous Ellsberg paradox (Ellsberg, 1961), which does not cause irrational behavior but can also not be explained by standard theory. In the Ellsberg paradox the decision maker is faced with choices that entail objective probabilities of events and choices where there is no objective probability attached to events, a phenomenon known as ambiguity. The component of standard expected utility theory that causes contradictions in this paradox is called the independence axiom. It says that if a choice is preferred to another choice, a mixture of the first and any other third choice should be preferred to a mixture with the second and this third choice. The in-

dependence axiom is an important cornerstone of the standard expected utility theory as it allows a simple representation by a utility function with a single belief about the world. To accommodate the Ellsberg paradox, the independence axiom needs to be weakened causing representations with (possibly) multiple beliefs about the world. Various weaker versions of the independence axiom have been suggested together with corresponding representation results, most notably maximin expected utility (Gilboa and Schmeidler, 1989) and choquet expected utility (Schmeidler, 1989). The contribution of this chapter is to provide a unified approach to these models by using a deep mathematical result on the existence of the Gâteaux derivative of the utility function almost everywhere. We provide a representation result, called local prior expected utility (LEU), under which the decision maker maximizes utility with respect to a set of priors. Under LEU the prior used to evaluate a certain act is equal to the gradient of some appropriately defined utility mapping. A closer look reveals that the interpretation of the gradient as a probabilistic belief is not purely technical but also natural. In our setting, the gradient expresses overall changes in utility due to small changes in the utility of a given state of the world. Thus, a strong increase in the overall utility due to a small change in the utility in a given state of the world means that the decision maker deems this state very likely. Using LEU as a starting point, one can simply impose further axioms, infer their impact on the utility function, and prove the corresponding representation results. All proofs in our framework are constructive, and therefore give a characterization of the sets of priors required for the axioms under investigation. The simplicity of the proofs and the characterization of the sets of priors distinguishes our approach from those presented in Schmeidler (1989), and Gilboa and Schmeidler (1989). Moreover, we show that prior representations can be obtained with even weaker axioms than Gilboa and Schmeidler's (1989) C-Independence axiom. The minimal independence axiom that still allows a prior representation is called Weak C-Independence, and preserves utility transformations with respect to translations. I believe this contribution will help future generations of decision and game theorist to easily understand decisions making under ambiguity and building new models more easily.

The chapters of this thesis can be read independently and contain detailed introductions and discussions. As each chapter discusses different issues, the notation will be introduced when needed and might differ across chapters.

Chapter 2

Beyond Coincidence: The Reasoning Process Underlying Utility Proportional Beliefs

When analyzing behavior in games we often restrict our attention solely to choices. Choice data not only drives the outcome in most settings but is also readily available. Nevertheless, we believe that secondary data can be a great source of insight when validating our models in experimental environments. In this chapter we show how experimental data on beliefs and reaction times can be used to increase our confidence in a particular model.¹

We focus on Bach and Perea's (2014), henceforth BP, concept of utility proportional beliefs. The main idea is to restrict players' beliefs in such a way that the differences in probabilities a player assigns to his opponent's choices should be equal to the differences in the opponent's utilities for these choices. Their concept is defined over beliefs, which can be advantageous as beliefs have a higher resolution than choices and therefore contain more information given the same number of observations. Our contribution is to provide a deeper understanding of their concept and to derive additional implications on peoples' behavior that can be tested. We derive an explicit expression to obtain utility proportional beliefs in two-player games without the need to apply BP's iterative algorithm. This

¹This chapter is based on my paper "Beyond coincidence: the reasoning process underlying utility proportional beliefs process."

		<i>Column Player</i>		
		<i>left</i>	<i>middle</i>	<i>right</i>
<i>Row Player</i>	<i>top</i>	5, 5	0, 4	2, 1
	<i>bottom</i>	4, 5	4, 4	1, 2

Figure 2.0.1: Example

direct approach can be interpreted as an iterative reasoning process and has implications for the players’ reaction times. The assumptions on reaction times are supported by both computational complexity and dual process theory. Using the data set on players’ beliefs from Nauerz et al. (2016)² we find evidence for the existence of utility proportional belief players and supporting evidence in their reaction times, strengthening our trust in the concept.

The concept of utility proportional beliefs is best illustrated by means of an example. Consider the game in Figure 2.0.1. Iterated elimination of strictly dominated strategies and therefore also Nash equilibrium predict that the *Column Player* will choose *left* and therefore the *Row Player* will play *top*. If the *Row Player* doubts, however, that the *Column Player* is fully rational and therefore also deems *middle* sufficiently likely, *bottom* seems to be a better choice. This way of thinking is quite common in real life situations. Utility proportional beliefs captures the idea of this tradeoff by the fact that the probabilities assigned to the choices have to be proportional to their utility. Following this idea, the *Row Player* would assign high probabilities to the *Column Player*’s choices *left* and *middle* because the difference in utilities for the two choices is rather small compared to the difference with the utility induced by *right*. Hence, it would be optimal for her to play *bottom*. Based on this simple idea, BP build a solution concept that does not rely on an equilibrium assumption and that is meant to take place in a single player’s mind.

For the two-player case, it turns out that the beliefs players can hold under the concept of utility proportional beliefs are unique, which allows us to find an explicit expression to compute players’ beliefs directly. The resulting expression can be interpreted as the players’ reasoning process under utility proportional beliefs. We find that the process resembles some intuitive properties. (1) Initial Belief: The players’ reasoning process starts with an initial belief about the op-

²Note that Chapter 3 is based on Nauerz et al. (2016).

ponent, which is the basis for the further belief formation process. It considers the goodness of the opponent's choices according to their average utility. (2) Interactive Reasoning Procedure: Within the reasoning process a player changes perspectives within every reasoning step. Hence, he considers his opponent's incentive structure and his own alternatingly. Every reasoning step starts with taking the opponent's perspective. First one forms a belief about one's own choice given the belief that one entertained about the opponent in the previous reasoning step. Then one takes the new belief of the opponent about oneself and forms a new belief about the opponent. The formation of the respective beliefs follows the same principle as the formation of the initial belief. Only that players do not consider the average utility but the expected utility given the opponent's belief from the previous reasoning step. This interpretation goes hand in hand with understanding the process as a series of reasoning steps that lead to a probabilistic assessment of another player's choices. It should not be understood as an exact description of every single computation that human beings perform but rather as a systematic process through which they analyze the game.

The remainder of this chapter is organized as follows. In Section 2.1 we introduce BP's epistemic definition of utility proportional beliefs and their algorithm that iteratively deletes those beliefs that cannot be hold under common belief in utility proportional beliefs. Then in Section 2.2, we develop an expression to obtain the players' beliefs directly, which can be interpreted as a reasoning process. Section 2.3 provides a cognitive interpretation and links it to dual processing theory. To illustrate the cognitive as well as the analytical features we discuss an example in Section 2.4. Finally, in Section 2.5 we present some experimental evidence supporting the existence of utility proportional belief players.

2.1 Epistemic Characterization

Consider a finite normal form game for two players, which is represented by a tuple

$$\Gamma = (I, (C_i)_{i \in I}, (U_i)_{i \in I})$$

where $I = \{1, 2\}$ denotes the set of players, C_i denotes player i 's finite choice set, and $U_i : \times_{j \in I} C_j \rightarrow \mathbb{R}$ denotes player i 's utility function. BP formalize utility proportional beliefs for normal form games using an epistemic model.

Definition 2.1. *An epistemic model of a game Γ is a tuple $\mathcal{M}^\Gamma = ((T_i)_{i \in I}, (b_i)_{i \in I})$,*

where

- T_i is a set of types for player $i \in I$
- $b_i: T_i \rightarrow \Delta(C_j \times T_j)$ assigns to every type $t_i \in T_i$ a probability measure with finite support³ on the set of opponents' choice-type combinations.

Epistemic types can be understood as mental states of a single person, where a type describes his beliefs about the opponent's choice-type combinations. The probability measure $b_i(t_i)$ represents type t_i 's belief function on the set of opponents' choice-type pairs. For convenience $b_i(t_i)$ also denotes any projected belief function for type t_i , e.g. $(b_i(t_i))(c_j|t_j)$ gives the probability that player i assigns to j 's choice c_j given that j is of type t_j . It will be clear from the context what $b_i(t_i)$ refers to. Now consider a game Γ , an epistemic model \mathcal{M}^Γ of it, and fix two players $i, j \in I$. A type $t_i \in T_i$ of i is said to *deem possible* some type $t_j \in T_j$ of her opponent j , if $b_i(t_i)$ assigns positive probability to an opponents' choice-type combination that includes t_j . The set $T_j(t_i)$ describes the set of types of player j deemed possible by t_i . Player i 's *conditional belief* that player j chooses c_j conditional on her type being t_j is expressed as $(b_i(t_i))(c_j|t_j) := (b_i(t_i))(c_j, t_j) / (b_i(t_i))(t_j)$ given a type $t_i \in T_i$ of player i , the opponent's type $t_j \in T_j(t_i)$, and her belief that j is of type t_j . Type t_i 's expected utility given her belief on her opponents' choice combinations is given by $u_i(c_i, t_i) = \sum_{c_j \in C_j} (b_i(t_i))(c_j) U_i(c_i, c_j)$. Additionally, type t_i 's average utility is denoted by $u_i^{\text{average}}(t_i) := |C_i|^{-1} \sum_{c_i \in C_i} u_i(c_i, t_i)$. Finally, $C := \times_{i \in I} C_i$ denotes all the choice combinations in the game. The best and the worst possible utilities of player i are denoted as $\bar{u}_i := \max_{c \in C} u_i(c)$ and $\underline{u}_i := \min_{c \in C} u_i(c)$. This concludes the formal representation of the epistemic model for utility proportional beliefs.

We are now ready to introduce the notion of expressing utility proportional beliefs.

Definition 2.2. Let $i, j \in I$ be the two players, and $\lambda_j \in \mathbb{R}$ such that $\lambda_j \geq 0$. A type $t_i \in T_i$ of player i expresses λ_j -utility-proportional-beliefs, if

$$(b_i(t_i))(c_j|t_j) - (b_i(t_i))(c'_j|t_j) = \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} (u_j(c_j, t_j) - u_j(c'_j, t_j)) \quad (2.1.1)$$

³The finite support condition restricts the focus on types that only assign positive probability to finitely many types for their opponent. This restriction avoids the introduction of a σ -algebra and therefore additional topological assumptions

for all $t_j \in T_j(t_i)$, for all $c_j, c'_j \in C_j$.

The definition directly corresponds to the idea of utility proportional beliefs: the difference in probabilities player i assigns to the opponents' choices is equal to the difference of the normalized utilities times the proportionality factor λ_j . This interpretation differs slightly from BP's in the sense that they call $\lambda_j / (\bar{u}_j - \underline{u}_j)$ the proportionality factor instead of λ_j . However, $1 / (\bar{u}_j - \underline{u}_j)$ normalizes the utilities such that maximal difference between the highest and lowest utility is equal to one and that beliefs are invariant with respect to affine transformations of the utilities (see BP). BP give an intuitive interpretation of the proportionality factor λ_j as measure of the sensitivity of a player's beliefs to differences in the opponents utilities. Note that there exists an upper bound for the proportionality factor called λ_j^{max} . It is the maximum value of λ_j for which equation (2.1.1) yields well-defined probability measures. The lower limit of λ_j is 0. In this case the differences in utilities do not play a role in the belief formation process. Players simply assign equal probability to their opponents' choices. BP also present an explicit formula for the belief about a given opponent's choice conditional of her being of a given type. It relates the conditional belief in a specific opponent's choice to the utility that this choice generates for the respective opponent.

Lemma 2.1. *Let $i, j \in I$ be the two players, and $\lambda_j \in \mathbb{R}_+$. A type $t_i \in T_i$ of player i expresses λ_j -utility-proportional-beliefs if and only if*

$$(b_i(t_i))(c_j|t_j) = \frac{1}{|C_j|} + \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} (u_j(c_j, t_j) - u_j^{average}(t_j)), \quad (2.1.2)$$

for all $t_j \in T_j(t_i)$, for all $c_j \in C_j$, and for $j \in I \setminus \{i\}$.

Equation (2.1.2) gives an intuitive interpretation of conditional beliefs under utility proportional beliefs: player i assigns to her opponent's type the uniform distribution on the respective opponent's choice set plus or minus an adjustment for each choice depending on its goodness relative to the average utility.

The concept of common belief in λ -utility-proportional-beliefs requires that both players entertain utility proportional beliefs, that both players believe their opponent holds utility proportional beliefs, that both players believe their opponents believe that their opponents do so, and so on. This idea is formally expressed by Definition 2.3.

Definition 2.3. *Let $i, j \in I$ be the two players, $t_i \in T_i$ be some type of player i , and*

$$\lambda = (\lambda_j)_{j \in I} \in \mathbb{R}_+^2.$$

- Type t_i expresses 1-fold belief in λ -utility-proportional-beliefs, if t_i expresses λ_j -utility-proportional-beliefs.
- Type t_i expresses k -fold belief in λ -utility-proportional-beliefs, if t_i only deems possible types $t_j \in T_j$ for j such that t_j expresses $k - 1$ -fold belief in λ -utility-proportional-beliefs, for all $k > 1$.
- Type t_i expresses common belief in λ -utility-proportional-beliefs, if t_i expresses k -fold belief in λ -utility-proportional-beliefs for all $k \geq 1$.

Definition 2.3 tells us that a type satisfying common belief in λ -utility-proportional-beliefs has λ_j -utility proportional-beliefs, believes that her opponent has λ -utility-proportional-beliefs, believes that her opponent believes that she holds λ -utility-proportional-beliefs, and so on. Finally, Definition 2.4 describes all choices player i can rationally make under common belief in λ -utility-proportional-beliefs.

Definition 2.4. Let $i, j \in I$ be the two players, and $\lambda = (\lambda_i)_{i \in I} \in \times_{i \in I} \mathbb{R}_+$. A choice $c_i \in C_i$ of player i is rational under common belief in λ -utility-proportional-beliefs, if there exists an epistemic model \mathcal{M}^Γ and some type $t_i \in T_i$ of player i such that c_i is optimal given $(b_i(t_i))$ and t_i expresses common belief in λ -utility-proportional-beliefs.

BP introduce an algorithm to find exactly those beliefs that are possible under common belief in λ -utility-proportional-beliefs. The algorithm iteratively deletes beliefs so that only beliefs, which are possible under common belief in λ -utility-proportional-beliefs, survive. To introduce the algorithm some further notation is required. By $P_i^0 := \Delta(C_j)$ BP denote the set of i 's beliefs about her opponent's choice combinations. Given $p_i \in P_i^0$ they define $u_i^{\text{average}}(p_i) := \sum_{c_i \in C_i} u_i(c_i, p_i) / |C_i|$. Furthermore, they define for player i and her opponent $j \in I \setminus \{i\}$, $p_i^* : \Delta(C_i) \rightarrow \Delta(C_j)$, where p_i^* is a function mapping the beliefs of player j on i 's choice to beliefs of i on j 's choice⁴:

$$(p_i^*(p_j))(c_j) := \frac{1}{|C_j|} + \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} (u_j(c_j, p_j) - u_j^{\text{average}}(p_j)) \quad (2.1.3)$$

for all $c_j \in C_j$ and for all $p_j \in P_j^0$. For a given $P_j \subseteq P_j^0$, BP denote by $p_i^*(P_j)$ the set $\{p_i^*(p_j) : p_j \in P_j\}$. These are all the ingredients needed to state the algorithm for iterated elimination of λ -utility-disproportional-beliefs formally.

⁴To save some notation we do not indicate that the function p_i^* also depends on λ_j . It will be clear from the context which λ_j we refer to.

Definition 2.5. For both players $i \in I$ and for all $k \geq 0$ the set P_i^k of i 's beliefs about her opponent's choices is inductively defined as follows:

$$\begin{aligned} P_i^0 &:= \Delta(C_j), \text{ and} \\ P_i^k &:= p_i^*(P_j^{k-1}). \end{aligned}$$

The set of beliefs $P_i^\infty = \bigcap_{k \geq 0} P_i^k$ contains the beliefs that survive iterated elimination of utility-disproportional-beliefs.

Using the algorithm repeatedly deletes beliefs that are not utility proportional with respect to beliefs from the preceding set of beliefs generated by the algorithm. Actually, it iteratively deletes beliefs p_i that cannot be obtained by the function p_i^* on the set P_j^{k-1} . Theorem 1 in BP establishes that this algorithm yields precisely those beliefs that a player can hold under common belief in λ -utility-proportional-beliefs. Furthermore, their Theorem 1 establishes that common belief in λ -utility-proportional-beliefs is always possible. Hence, the concept can be applied to describe the players' reasoning in any static game.

2.2 Computing Beliefs Directly

As shown above, one can find a player's belief expressing common belief in λ -utility-proportional-beliefs by applying BP's algorithm that iteratively eliminates utility-disproportional beliefs. This iterative procedure, however, allows only for limited insight into the players' reasoning process and is more focused on the outcome than the process itself. In this section we provide a different view using a direct approach to obtain the players' beliefs under common belief in λ -utility-proportional-beliefs, allowing us more insight into the belief formation process.

Proposition 2.1 establishes that there is a unique fixed point of the $p_i^* \circ p_j^*$ mapping. The fixed point is equal to the belief that player i can hold about player j under common belief in λ -utility-proportional-beliefs.

Proposition 2.1. *There exists a unique fixed point for $p_i^* \circ p_j^*$.*

Proof. See Appendix 2.A. □

In the following we will show that there exists an expression to compute the fixed point directly and therefore the belief that a player can hold under common

belief in λ -utility-proportional-beliefs. Before we can proceed we need to introduce some additional notation. We will denote $((p_i^*(p_j))(c_j))_{c_j \in C_j}$ by the vector $p_i^*(p_j)$. For every $n \times m$ game Γ we denote the number of choices of player i by $n = |C_i|$ and the number of choices for player j by $m = |C_j|$. The $n \times 1$ vector i_n with $i_n = (\frac{1}{n}, \dots, \frac{1}{n})$ is the vector equivalent to $1/|C_i|$ in (2.1.3). Let $C_i = \{c_i^1, \dots, c_i^n\}$ and $C_j = \{c_j^1, \dots, c_j^m\}$ so that we can denote player i 's $n \times m$ utility matrix by

$$U_i = \begin{bmatrix} U_i(c_i^1, c_j^1) & \cdots & U_i(c_i^1, c_j^m) \\ \vdots & & \vdots \\ U_i(c_i^n, c_j^1) & \cdots & U_i(c_i^n, c_j^m) \end{bmatrix}. \quad (2.2.1)$$

Now we can express the expected utility as a $n \times 1$ vector $u_i(p_i) = U_i p_i$, the average utility as a $n \times 1$ vector $u_i^{average} = \frac{1}{n} \mathbb{1}_n u_i(p_i)$ with every element equal to the average utility, where $\mathbb{1}_n$ represents a $n \times n$ matrix of 1's. After substituting $u_i(p_i)$ into $u_i^{average}$ and rearranging we obtain

$$u_i(p_i) - u_i^{average} = (I_n - \frac{1}{n} \mathbb{1}_n) U_i p_i,$$

where I_n is the $n \times n$ identity matrix. The centering term $(I_n - \frac{1}{n} \mathbb{1}_n)$ will play a crucial role. Hence, we will denote the term by the $n \times n$ centering matrix $C_n := (I_n - \frac{1}{n} \mathbb{1}_n)$. The matrix C_n has $\frac{n-1}{n}$ on the diagonal and $-\frac{1}{n}$ off the diagonal. Intuitively, the centering matrix subtracts the mean from the columns of a matrix. Now can write $p_i^*(p_j)$ as

$$p_i^*(p_j) = i_m + \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} C_m U_j p_j. \quad (2.2.2)$$

Since $p_i^*(p_j)$ maps j 's beliefs about i into i 's beliefs about j , and $p_j^*(p_i)$ maps i 's beliefs about j into j 's beliefs about i , we can substitute $p_j^*(p_i)$ for p_j , so that we obtain $p_i^*(p_j^*(p_i))$. Suppose $p'_i \in P_i^\infty$ is the fixed point of $p_i^* \circ p_j^*$ then we can solve $p'_i = p_i^*(p_j^*(p'_i))$ for p'_i and obtain i 's beliefs about j under common belief in λ -utility-proportional-beliefs. We substitute p_j by player j 's belief generating

function in (2.2.2) and obtain

$$\begin{aligned} p'_i = p_i^*(p_j^*(p'_i)) &= i_m + \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} C_m U_j p_j^*(p'_i) \\ &= i_m + \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} C_m U_j \left[i_n + \frac{\lambda_i}{\bar{u}_i - \underline{u}_i} C_n U_i p'_i \right]. \end{aligned}$$

Solving for p'_i yields

$$p'_i = (I_m - \frac{\lambda_j \lambda_i}{(\bar{u}_j - \underline{u}_j)(\bar{u}_i - \underline{u}_i)} C_m U_j C_n U_i)^{-1} (i_m + \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} C_m U_j i_n). \quad (2.2.3)$$

Hence, we can express player i 's beliefs about player j as an explicit formula.⁵ The proportionality factor λ_j affects the extent to which player i takes j 's differences in utility into account when reasoning about player j . To have a standardized measure, we define $\alpha_i := \lambda_i / \lambda_i^{max}$ because we need $\lambda_i < \lambda_i^{max}$ to hold and λ_i^{max} differs from game to game. Remember λ_i^{max} is defined so that $p_j^*(p_i)$ always yields well-defined probability measures. Consequently, we write $\alpha_i \lambda_i^{max}$ instead of λ_i with $\alpha_i \in [0, 1)$ to have a standardized measure for the players' sensitivity to differences in the utilities of their respective opponents. To save some further notation we define $U_i^{norm} = (\lambda_i^{max} / (\bar{u}_i - \underline{u}_i)) U_i$. Moreover, we define the matrix $G_j := \alpha_j C_m U_j^{norm}$. Since this matrix appears several times in the expression above, it is useful to develop a more intuitive understanding. By left-multiplying the normalized utility matrix U_j^{norm} with the centering matrix C_m , one obtains a matrix where for every element the average of its column has been subtracted. Note that the rows of U_j^{norm} correspond to i 's choices and the columns to j 's choices. The same holds for the matrix $C_m U_j^{norm}$, only that now each element represents the relative goodness of a choice given an opponent's choice. Therefore, the matrix G_j gives the goodness of a choice given an opponent's choice, scaled by the sensitivity to the opponents differences in utility α_j . Now we can define $\beta_i(G_i, G_j)$ as the function for i 's belief about j under common belief in λ -utility-proportional-beliefs as

$$\beta_i(G_i, G_j) := (I_m - G_j G_i)^{-1} (i_m + G_j i_n). \quad (2.2.4)$$

⁵The existence of the inverse is discussed in remark 6.

With $\beta_i(G_j, G_i)$ we obtain i 's beliefs about j directly given their respective payoff matrices and proportionality factors. This expression, however, involves the inverse of a matrix which makes the interpretation quite difficult. To obtain an expression for the players' beliefs that does not rely on the inverse of a matrix, note that the inverse can be expressed as an infinite sum. Thus, we define the $m \times m$ matrix $S_j := G_j G_i$ and express the inverse in (2.2.4) as $(I_m - S_j)^{-1}$. To show that $(I_m - S_j)^{-1} = I_m + S_j + S_j^2 + S_j^3 + \dots$ holds, we first have to convince ourselves that the infinite sum of matrices always converges.

Lemma 2.2. *The infinite sum $(I_m + \sum_{n=1}^{\infty} S_j^n)$ always converges to $(I_m - S_j)^{-1}$.⁶*

Proof. See Appendix 2.A. □

Now we can express the players' beliefs under common belief in λ -utility-proportional-beliefs directly using the infinite sum

$$\begin{aligned} \beta_i &= \sum_{k=0}^{\infty} (G_j G_i)^k (i_m + G_j i_n) & (2.2.5) \\ &= (i_m + G_j i_n) + G_j G_i (i_m + G_j i_n) \\ &\quad + G_j G_i [G_j G_i (i_m + G_j i_n)] + \dots \end{aligned}$$

2.3 Cognitive Characterization

Using equation (2.2.5) we can directly compute the beliefs that players hold under common belief in λ -utility-proportional-beliefs. This expression can be interpreted as a reasoning process. In the following we will discuss how the mathematical properties could be linked to such a cognitive process.

First observe that the expression $(i_m + G_j i_n)$ is repeated several times and that $(i_m + G_j i_n)$ is (2.2.2) applied to the uniform belief about player i 's choices. In the second term of equation (2.2.5), this expression is then adjusted by left multiplying the matrices $G_j G_i$. In the third term the second term is adjusted by left multiplying $G_j G_i$, and so on. Therefore, we call $(i_m + G_j i_n)$ the initial belief, $\beta_i^{initial}$. To further emphasize this process, we define β_i^k as the belief that player i

⁶Since the infinite sum $(I_m + \sum_{n=1}^{\infty} S_j^n)$ converges, the inverse $(I_m - S_j)^{-1}$ is guaranteed to exist.

holds after the k th reasoning step,

$$\begin{aligned}\beta_i^0 &:= \beta_i^{initial} \\ \beta_i^k &:= \beta_i^{initial} + G_j G_i (\beta_i^{k-1}),\end{aligned}\tag{2.3.1}$$

such that $\lim_{k \rightarrow \infty} \beta_i^k = \beta_i$ holds.

The initial belief $\beta_i^{initial}$ is the core of the step-wise adjustment process described in (2.3.1). Moreover, it shows how player i constructs her beliefs about player j without taking into account that player j reasons about her. This is the same as when player i would believe that player j is insensitive to her differences in utilities in the sense that $\alpha_i = 0$, so that $\beta_i = \beta_i^{initial}$. In this case, player i 's belief about player j only depends on j 's payoff matrix and not on her own. Consequently, player i does not consider any strategic interaction when reasoning about j . Let us now consider the simplest case when player i is insensitive to differences in j 's utilities, i.e. $\alpha_j = 0$. In this case, we have $\beta_i^{initial} = i_m$. Recall that the $m \times 1$ vector $i_m = (\frac{1}{m}, \dots, \frac{1}{m})$ assigns a uniform distribution over j 's choices. Hence, if $\alpha_j = 0$ player i deems all choices of player j equally likely. If we have $\alpha_j > 0$ player i takes j 's payoff scheme into account when forming her beliefs. As in the previous case player i starts off by assigning equal probability to her opponent's choices. In this case, however, the uniform assignment will be corrected by adding the term $G_j i_n$. As discussed above, G_j gives the relative goodness of j 's choices given i 's choice. We obtain the goodness of j 's choices given that player j assigns equal probability to i 's choices. Therefore, when player i forms her belief about player j she assumes that j deems each of her choices equally likely when assessing the goodness of j 's choices given i 's. By adding up the vectors $G_j i_n$ and i_m we obtain the belief that player i holds about j when we only consider the initial belief. Intuitively, player i first assigns equal probability to all of j 's choices and then corrects these probabilities according to the average goodness of j 's choices.

Every following step in the reasoning process makes adjustments to the previous belief, as can be seen in (2.3.1). The process starts with $\beta_i^{initial}$, which is the basis for further reasoning and therefore not considered a reasoning step itself. In the first reasoning step player i adjusts the initial belief by adding the initial belief left multiplied by $G_j G_i$. To obtain a more intuitive understanding

we rewrite β_i^1 as follows

$$\beta_i^1 = i_m + G_j(i_n + G_i(\beta_i^0)).$$

We see that first player i takes j 's perspective, which is reflected in the expression $i_n + G_i(\beta_i^0)$. Here player j forms a belief about player i given i 's belief about j from the previous reasoning step, the initial belief. First j assigns equal probability to all of i 's choices. Then she corrects these beliefs by the goodness of i 's choices given i 's belief about j from the previous reasoning step. The result is a new belief of j about i . Then player i takes her own perspective and assigns equal probability to all of j 's choices. These probabilities are then again corrected by the goodness of j 's choices given the new belief of j about i . In the second reasoning step, we have a very similar procedure. The second reasoning step can best be expressed as

$$\beta_i^2 = i_m + G_j(i_n + G_i(\beta_i^1)).$$

As one can immediately see, the second step follows exactly the same process as the first. The basis for the reasoning step is i 's belief about j from the previous reasoning step. Based on this belief, i takes j 's perspective and forms a belief about herself. This belief is then used to form a new belief about j , resulting in the final belief that i holds about j after the second reasoning step. This exact procedure is then being repeated for every single reasoning step.

It is also important to note that later reasoning steps will be less important for the final belief than earlier ones. Note that $G_j = \alpha_j \lambda_j^{max} C_m U_j^{norm}$, so that (2.2.5) can be written as

$$\beta_i = \sum_{k=0}^{\infty} (\alpha_i \alpha_j \lambda_i^{max} \lambda_j^{max} C_m U_j^{norm} C_n U_i^{norm})^k \beta_i^{initial}.$$

Since $\alpha_i, \alpha_j \in [0, 1)$, later terms in $\sum_{k=0}^{\infty} (\alpha_i \alpha_j \lambda_i^{max} \lambda_j^{max} C_m U_j C_n U_i)^k$ will be smaller than earlier ones and therefore less important for the final belief β_i .⁷ This has also an important implication for the meaning of the proportionality factor α_i : the lower the value of α_i the fewer steps of reasoning a player will undergo to approximate the final belief within a reasonable bound. The same holds true for her opponent's proportionality factor.

These features correspond closely to findings in the psychology literature.

⁷This statement is valid since S_j^k converges to the zero matrix as $k \rightarrow \infty$ for all $\alpha_i, \alpha_j \in [0, 1)$.

System 1	System 2
associative	rule-based
holistic	analytic
relatively undemanding of cognitive capacity	demanding of cognitive capacity
relatively fast	relatively slow
acquisition by biology, exposure, and personal experience	acquisition by cultural and formal tuition

Table 2.1: Properties of System 1 & 2

In his book “Thinking, Fast and Slow” Kahneman (2011) advocates the idea of reasoning in two distinct ways. He calls the two modes of thinking System 1 and System 2, according to Stanovich et al. (2000). Note that the word system should not indicate an actual system but only serves as label for different modes of thinking. Stanovich et al. (2000) describe System 1 as an automatic and mostly unconscious way of thinking that demands little computational capacity. System 2 describes the idea of deliberate reasoning. It comes into play when controlled analytical thinking is needed. Table 2.1 summarizes the properties of the two systems according to Stanovich et al. (2000).

The idea of System 1 describes the unconscious first reaction to a situation, which happens almost immediately and without demanding a lot of cognitive energy. Moreover, Kahneman (2011) argues that the beliefs formed by System 1 are the basis for conscious reasoning within System 2. This is consistent with the initial belief. As we showed, the initial belief does not take into account any strategic interaction. It can be seen as an automatic initial reaction to the game. The reasoning process, which occurs after the initial belief is formed, can be imagined as being executed by System 2. Taking another player’s perspective takes deliberate reasoning and can hardly be done automatically. Think of dominance solvable games, one has an intuitive impression of how the opponent might choose but to check whether that intuition is correct, requires the effort to repeatedly take the other player’s perspective and check for her reactions.

		<i>Column Player</i>	
		<i>left</i>	<i>right</i>
<i>Row Player</i>	<i>top</i>	320, 40	40, 80
	<i>bottom</i>	40, 80	80, 40

Figure 2.4.1: Asymmetric matching pennies

2.4 Example

The game in Figure 2.4.1 is taken from Goeree and Holt (2001) and also serves as an example in BP. In this game utility proportional beliefs predicts that the *Row Player* believes that the *Column Player* chooses *left* with a probability of 0.37 and *right* with a probability of 0.63. For the *Column Player* it predicts that she believes that the *Row Player* will choose *top* with a probability of 0.63 and *bottom* with a probability of 0.37 when taking both sensitivity parameters to be maximal.⁸

We will use this example to illustrate details of the reasoning process and to provide additional analytical insights that follow from our approach. To obtain an understanding of the players' reasoning process in this game, we first construct their goodness of choice matrices G_R for the *Row Player* and G_C for the *Column Player*. Making the players as sensitive as possible to differences in pay-offs will help to illustrate the process's properties. Hence, we set α_R and α_C close to 1. For the *Row Player* we have

$$G_R = \alpha_R \frac{\lambda_R^{max}}{\bar{u}_R - \underline{u}_R} C_2 U_R = \begin{bmatrix} 0.5 & -0.07 \\ -0.5 & 0.07 \end{bmatrix}$$

and for the *Column Player*

$$G_C = \alpha_C \frac{\lambda_C^{max}}{\bar{u}_C - \underline{u}_C} C_2 U_C = \begin{bmatrix} -0.5 & 0.5 \\ 0.5 & -0.5 \end{bmatrix}.$$

For convenience all presented numbers are rounded. Let us now examine how the *Row Player* forms her beliefs about the *Column Player*.

Row Player: The initial belief assigns equal probability to both of the *Column Player's* choices. The *Row Player* believes initially that the *Column Player* believes

⁸For this game we have $\lambda_i^{max} = \lambda_j^{max} = 1$.

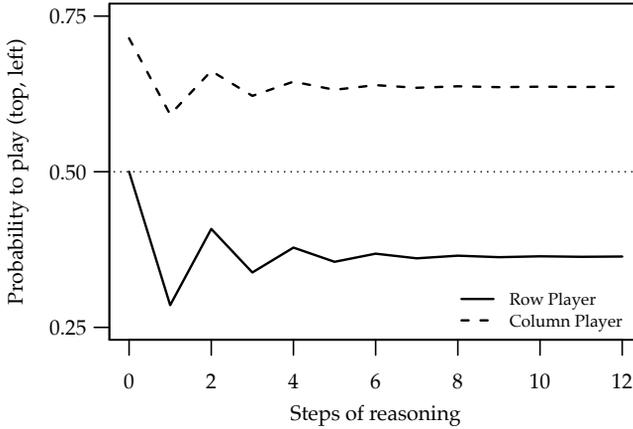


Figure 2.4.2: Players' belief formation process

that she chooses *top* and *bottom* with equal probability. But when the *Column Player* deems both choices equally likely her expected goodness will be 0 for both of her choices. As the goodness of the *Columns Player's* choices alone does not yield any additional information, the *Row Player* deems both of her choices equally likely in her initial belief. Note that it is just a coincidence that the *Row Player's* initial belief deems both of the *Column Player's* choices equally likely, which is due to the fact that the *Column Player* is indifferent between her choices when she believes that the *Row Player* play both choices with equal probability.

Figure 2.4.2 shows how the *Row Player's* belief develops with each reasoning step. After the seventh reasoning step the belief stabilizes and changes only marginally in the following reasoning steps. The unique belief of the *Row Player* under common belief in λ -utility-proportional-beliefs is $\begin{bmatrix} 0.37 & 0.63 \end{bmatrix}$. Consequently, the *Row Player* deems the *Column Player's* choice *right* more likely than *left*.

Column Player: Again, we first consider the initial belief of the *Column Player* about the *Row Player*. As the *Row Player's* payoffs are quite distinct the *Column Player* deems the *Row Player's* choice *top* significantly more likely than her choice *bottom*, we have $\beta_C^{initial} = \begin{bmatrix} 0.71 & 0.29 \end{bmatrix}^t$. This outcome is quite intuitive when the *Column Player* believes that the *Row Player* does not consider any strategic interaction and deems the *Column Player's* choices equally likely.

Also the *Column Player's* beliefs converge quickly (see Figure 2.4.2). Interestingly, the *Column Player's* beliefs do not fluctuate as strongly as the *Row Player's*

and converge quicker. The resulting unique belief under common belief in λ -utility-proportional-beliefs is $[0.63 \ 0.37]$. Hence, the *Column Player* deems the *Row Player's* choice *top* more likely than her choice *bottom*.

Finally, we want to point out some analytical properties resulting from the direct approach. In Section 2.2 we showed that beliefs players hold under common belief in utility proportional beliefs can be obtained directly via the fixed point of $p_i^* \circ p_j^*$. Figure 2.4.3 depicts this situation for our example. The point where the two lines corresponding to p_i^* and p_j^* cross indicates the fixed point of the mapping and thus the beliefs that players can hold under common belief in utility proportional beliefs. Therefore, the players' beliefs in this example can be simply obtained by finding the intersection of the two players' utility proportional responses p_i^* and p_j^* . Furthermore, note that as we reduce α_i the line corresponding to p_i^* will become steeper, where the dot on the line indicates its turning point as we vary α_i . The line corresponding to p_j^* becomes flatter as we decrease α_j but goes through the point on the line for all values of α_j . We can use this insight to derive the set of all possible predictions of utility proportional beliefs as we vary the α_i and α_j . The resulting set is indicated by the gray area in Figure 2.4.3. Furthermore, if we were curious how the solution changes when increasing the *Row Player's* payoff for *top* and *left*, we could simply take the derivative with respect to that payoff and infer that the slope of the line corresponding to p_j^* would decrease. In this case, however, it would turn around the point $(0, 1)$, increasing the set of possible beliefs.

2.5 Experimental Evidence

In this section we present evidence that utility proportional belief players exist and provide additional support from a cognitive perspective using the characterization of utility proportional beliefs as a reasoning process. The data set we use stems from Nauerz et al. (2016), who test whether people's beliefs are consistent with a given reasoning concept across a number of games. Their experiment was implemented in an anonymized laboratory setting, at the Behavioral and Experimental Economics Laboratory at Maastricht University, with a total of 88 subjects. Subjects stated beliefs in 24 different bi-matrix games, for which they were randomly matched into pairs. For each game, subjects were both asked to indicate their belief about their opponent's choices and make their own de-

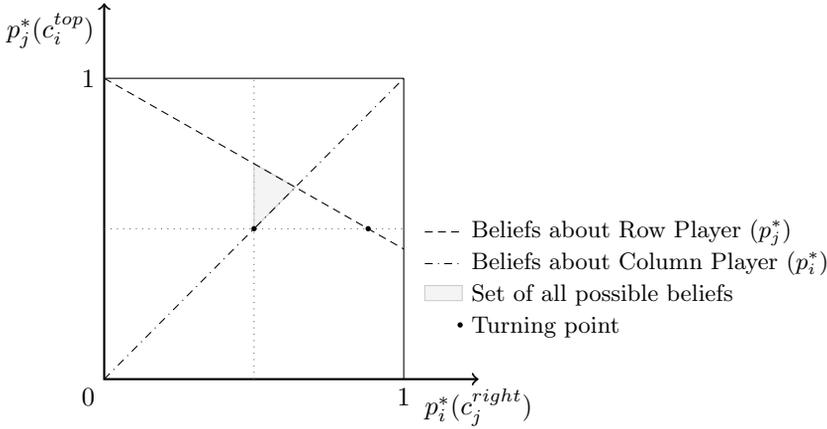


Figure 2.4.3: Players' utility proportional responses

cision. To elicit their beliefs, they were asked about the probability they assign to one of their opponent's choices. These statements were incentivized using the L-adjusted quadratic scoring rule (Offerman and Palley, 2015) in order to correct for risk and loss aversion. Subjects did not receive any feedback about their opponents' beliefs or choices throughout the experiment in order to avoid learning. At the end of the experiment, subjects were informed about their payoff, which was determined by either one of the subject's decisions or one of his stated beliefs at random in order to mitigate hedging effects (Blanco et al., 2010). The average payoff for a participants was 30 EUR for a session of two hours.⁹

Nauerz et al. (2016) designed the games such that deeming one of the opponent's choices more likely than another can only be explained by one solution concept. Hence, they are able to attribute solution concepts to subjects without having to estimate a statistical model. An example of one of their games will follow shortly. Nauerz et al.'s (2016) assume that $\alpha_i > 0$ and $\alpha_j > 0$ in order to make sharp distinctions between the tested concepts. This restriction is not needed in our analysis and therefore dropped. In the following, we exclusively consider subjects who stated beliefs in line with utility proportional beliefs in all 24 games and use the specific design of the games to distinguish beliefs. Since the players' most extreme beliefs are not necessarily obtained when both sensitivity parameters are taken to be maximal, as can be seen in Figure 2.4.3, we find the players' maximal beliefs under utility proportional beliefs via non-linear op-

⁹For more details about the experiment, consult Nauerz et al. (2016).

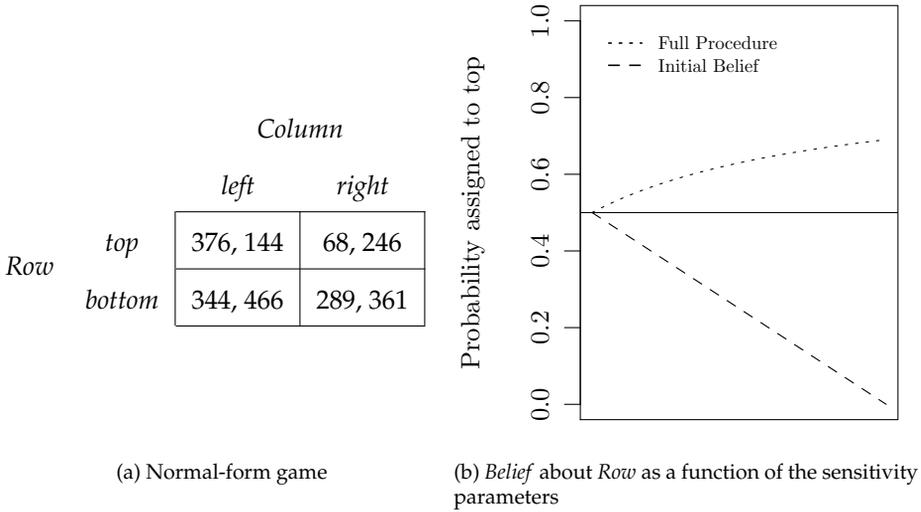


Figure 2.5.1: Example of game to distinguish level-2 beliefs

timization under the restriction that all steps in the reasoning process are well defined. We find that 6 of the 88 subjects hold utility proportional beliefs in all games—evidence of the existence of utility proportional beliefs.

In principle the result above could still be attributed to chance as we only test if subjects’ beliefs coincide with the predictions of utility proportional beliefs. Using insights from the cognitive characterization of utility proportional beliefs in Section 2.3 we develop additional hypotheses on the players’ reaction times to provide further evidence in favor of the existence of utility proportional belief players. In the simplest version of the reasoning process (i) a player neither takes his own nor his opponent’s incentives into account when forming his belief about the opponent ($\alpha_i = \alpha_j = 0$). Then his belief is simply a uniform distribution over the opponent’s choices. Therefore, we hypothesize that his belief formation should be almost instantaneous. The next step in the reasoning process is (ii) to form an initial belief. When forming the initial belief a player believes that his opponent believes that he himself plays uniformly and uses the resulting incentives of his opponent to form his belief ($\alpha_i = 0$ and $\alpha_j > 0$). The initial belief then serves as a basis for further reasoning where the player iteratively considers his and his opponent’s incentives. Thus, the initial belief is formed earlier in the reasoning process and therefore it is reasonable to assume that such a belief needs

less processing time than a belief formed under the (iii) full procedure (where $\alpha_i > 0$ and $\alpha_j > 0$). Hence, we hypothesize that a player who believes that his opponent does not take his incentives into account, and therefore only forms an initial belief, will be quicker when stating his beliefs than a player who believes that his opponent takes his incentives into account.

Figure 2.5.1 serves as an example of the types of games used in Nauerz et al. (2016). The game is designed to distinguish certain kinds of beliefs: in case (i) the *Column Player* deems both of his opponent's choices equally likely; under the assumptions of case (ii) the *Column Player* has to deem *bottom* more likely than *top*; and finally in case (iii) he would have to deem *top* more likely than *bottom*. Therefore, Nauerz et al. (2016) experimental setup provides the data we need to test our hypotheses about the respective reaction times in each of the three cases.

Table 2.2: Differences in Reaction Times

	<i>Dependent variable:</i>	
	Reaction Time Belief	
	(a) All Periods ¹⁰	(b) From period 8
Initial Belief (i)	-15.538*** (3.980)	-13.621*** (3.696)
Full Procedure (iii)	2.303 (4.620)	1.175 (4.107)
Constant	74.462*** (7.015)	26.545*** (3.422)
Observations	144	102
R ²	0.519	0.276
Adjusted R ²	0.396	0.262
Residual Std. Error	14.020 (df = 114)	11.350 (df = 99)
F Statistic	4.238*** (df = 29; 114)	18.883*** (df = 2; 99)
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

To test our hypotheses we ran regression analyses with "Reaction Time Beliefs" as dependent variable and dummy variables for the three different kinds of beliefs as independent variables. The initial belief (ii) was omitted and serves as a benchmark. Table 2.2 shows the results of two models. In model (a) we included all periods of the experiment and dummies for every period. Since

¹⁰This model also includes dummies for all 24 periods, which are all significant but not shown for a matter of presentation.

subjects are significantly slower in the first periods, we ran a second model (b) to exclude possible effects of learning¹¹ on the reaction time where we excluded the first 8 periods. In both models the coefficient for uniform beliefs is lower than zero, and the coefficient for the full procedure is larger than the zero.¹² Therefore, our results suggest that the time a player takes to state his belief is linked to the type of belief he states, which is in line with our interpretation of the reasoning process under utility proportional beliefs.

In conclusion, we not only find that 6 of 88 subjects coincide with the predictions of utility proportional beliefs in all games but also that the time they take to form their beliefs is in line with our cognitive interpretation of utility proportional beliefs.

2.6 Conclusion

Traditionally, game theoretic concepts have been judged by how well they fit a given data set. In this chapter we provide conditions for subjects' reaction times under the hypothesis that they are utility proportional belief types.

Instead of using BP's iterative algorithm we find an explicit expression to obtain the players' beliefs under utility proportional beliefs. This expression can be interpreted as the players' reasoning process, yielding conditions for the players' reaction times. Experimental evidence from Nauerz et al.'s (2016) data set supports the existence of utility proportional belief players. Moreover, these players confirm to our cognitive conditions, increasing our trust in the concept.

¹¹Here we mean learning effects in the sense of becoming more familiar with stating beliefs in strategic situations. There was no feedback that would have allowed actual learning to take place.

¹²The coefficients for (i) are in both models significantly different from 0 on a 5 percent level.

2.A Appendix - Proofs

Proof of Proposition 2.1. First, remember that P_i^k denotes the set of beliefs generated for player i in round k of the algorithm. Define for any two sets $A, B \subseteq P_i^0$ and for all $\alpha \in [0, 1]$ the set

$$\alpha A + (1 - \alpha)B := \{\alpha a + (1 - \alpha)b : a \in A \text{ and } b \in B\}.$$

Let $p_i^*(\cdot, \lambda_j)$ denote the function p_i^* induced by the proportionality factor λ_j . BP's Lemma 4 states that for the k^{th} iteration of the algorithm there exists some belief $p_i \in P_i^0$ such that the convex combination $\alpha^k P_i^0 + (1 - \alpha^k)\{p_i\}$ contains all beliefs that a player can hold under up to k -fold belief in λ -utility-proportional-beliefs. For the case of common belief in λ -utility-proportional-beliefs, we let k go to infinity and obtain

$$\lim_{k \rightarrow \infty} \alpha^k P_i^0 + (1 - \alpha^k)\{p_i^k\} = \{p_i^\infty\} \quad (2.A.1)$$

for some $p_i^\infty \in P_i^0$ where p_i^k denotes a belief $p_i^k \in P_i^0$ such that $P_i^k \subseteq \alpha^k P_i^0 + (1 - \alpha^k)\{p_i^k\}$ holds. We observe that the convex combination converges to a singleton and hence to a unique fixed point. Now suppose $p'_i \in P_i^\infty$ is the belief of i and $p'_j \in P_j^\infty$ is the belief of j for which equation (2.A.1) holds. Recall, that by the construction of the algorithm it holds that $P_i^k = p_i^*(P_j^{k-1}, \lambda_j)$. Then it holds that $p_i^*(p'_j) = p'_i$ and $p_j^*(p'_i) = p'_j$ and hence it also holds that $p_i^*(p_j^*(p'_i)) = p'_i$. Consequently, the mapping $p_i^* \circ p_j^*$ has a unique fixed point in two-player games where the players' are holding common belief in λ -utility-proportional-beliefs. \square

Lemma 2.3. *It holds that $(p_i^* \circ p_j^*)^k(p_i) = (i_m + G_j i_n) + S_j(i_m + G_j i_n) + \dots + S_j^{k-1}(i_m + G_j i_n) + S_j^k p_i$ for all $p_i \in P_i^0$ for all $k = \{1, \dots, \infty\}$.*

Proof. We prove Lemma 2.3 by induction on k . Algebraically, the composed function $p_i^* \circ p_j^*$ corresponds to

$$\begin{aligned} p_i^*(p_j^*(p_i)) &= i_m + G_j(i_n + G_i p_i) \\ &= (i_m + G_j i_n) + S_j p_i. \end{aligned}$$

Hence, we have for $k = 1$ that $(p_i^* \circ p_j^*)(p_i) = (i_m + G_j i_n) + S_j p_i$.

For some $k > 1$, suppose the $k - 1$ previous steps were iteratively constructed

by substituting $p_i^* \circ p_j^*$ for p_i at every step. Then we have

$$(p_i^* \circ p_j^*)^{k-1}(p_i) = (i_m + G_j i_n) + S_j(i_m + G_j i_n) + S_j^2(i_m + G_j i_n) + \cdots + S_j^{k-1} p_i.$$

Now substitute p_i by $(p_i^* \circ p_j^*)(p_i)$ to obtain step k

$$\begin{aligned} (p_i^* \circ p_j^*)^k(p_i) &= (i_m + G_j i_n) + S_j(i_m + G_j i_n) + \cdots + S_j^{k-1} ((i_m + G_j i_n) + S_j p_i) \\ &= (i_m + G_j i_n) + S_j(i_m + G_j i_n) + \cdots + S_j^{k-1}(i_m + G_j i_n) + S_j^k p_i, \end{aligned}$$

which is what we wanted to prove. \square

Proof of Lemma 2.2. According to Lemma 2.3 we have that

$$(p_i^* \circ p_j^*)^k(p_i) = (i_m + G_j i_n) + S_j(i_m + G_j i_n) + \cdots + S_j^{k-1}(i_m + G_j i_n) + S_j^k p_i,$$

for every $k \in \mathbb{N}$ and some $p_i \in P_i^0$. Factoring out $(i_m + G_j i_n)$ yields

$$(p_i^* \circ p_j^*)^k(p_i) = (I_m + S_j + \cdots + S_j^{k-1})(i_m + G_j i_n) + S_j^k p_i.$$

Since BP showed in their Theorem 1 that the iterative application of p_i^* and p_j^* yields exactly those beliefs that the players can hold under common belief in λ -utility-proportional-beliefs, and in their Theorem 2 that these beliefs must be unique in the two player case, $(p_i^* \circ p_j^*)^k(p_i)$ must also converge to the unique belief for $k \rightarrow \infty$. We then have $\lim_{k \rightarrow \infty} (p_i^* \circ p_j^*)^k(p_i) = (I_m + \sum_{n=1}^{\infty} S_j^n)(i_m + G_j i_n)$, which shows that the infinite sum $(I_m + \sum_{n=1}^{\infty} S_j^n)$ exists. It also follows that S_j^k converges to the zero matrix for $k \rightarrow \infty$ since $(p_i^* \circ p_j^*)^k(p_i)$ converges to the unique belief for $k \rightarrow \infty$ and since S_j^k is the only element affected by k . As the infinite sum $(I_m + \sum_{n=1}^{\infty} S_j^n)$ exists and S_j^k converges to the zero matrix, it follows that the infinite sum $(I_m + \sum_{n=1}^{\infty} S_j^n)$ always exists and that it converges to $(I_m + S_j)^{-1}$. \square

Chapter 3

Identifying Strategic Thinking in One-Shot Games

Strategic interactions with perfectly rational agents are well understood. Experimental evidence, however, suggests that human beings do not behave in a perfectly rational way, especially in one-shot settings. First, when facing a new strategic decision problem without prior knowledge, it might not be reasonable to assume that the opponents are perfectly rational. Second, even if an agent would assume that his opponents are perfectly rational, he might not be willing or able to perform computations with the accuracy needed to act in a perfectly rational way. Understanding one-shot strategic interactions is important as many economic interactions only occur once, e.g. in negotiations between unacquainted buyers and sellers. Even for repeated interactions the initial response might play a crucial role as it can affect the eventual equilibrium. Thus it is important to understand how agents¹ approach these strategic decisions.²

Experiments offer the possibility to acquire information about subjects' initial response to strategic interactions which is often not available in real world settings. They allow researchers to augment evidence from choice data to increase confidence in models beyond their fit to the data. In this paper, we use stated beliefs about the behavior of their opponent to explore whether subjects'

¹We use the words agent and player to describe individuals theoretically. Agent is used within a more general economic environment and player if we specifically talk about games.

²This chapter is based on the paper "Identifying Strategic Thinking in One-Shot Games" that I wrote together with Frauke Meyer and Marion Collewet.

strategic thinking can be explained by different behavioral solution concepts, namely level-k (Lk) (Stahl and Wilson, 1994), quantal response equilibrium (QRE) (McKelvey and Palfrey, 1995), and utility proportional beliefs (UPB) (Bach and Perea, 2014). For this purpose, we elicited subjects' beliefs, as well as their choices and the respective time they needed to perform each task in a laboratory experiment. Subjects stated probabilistic beliefs and made choices in 28 one-shot 2×2 and 2×3 normal-form bimatrix games without feedback in between games. For each game, subjects first stated a probabilistic belief about their opponent's choice and then made their own decision. Beliefs were elicited by means of the L-adjusted Quadratic Scoring Rule (Offerman and Palley, 2015), which accounts both for risk and loss aversion.

Using beliefs instead of choices to infer subjects' approach to games provides us with more informative data. Consider the case where a belief induces almost equal expected utility of two choices. Unless an agent is perfectly rational, i.e., performs computations with perfect accuracy, a choice might reflect his assessment or might simply be a mistake. In economics we often allow for this possibility by using the logit choice function, which assumes that bigger mistakes are less likely than small mistakes when reacting to a belief. Therefore, some choices certainly carry important information whereas others might contain little information—a problem that is less prominent when considering beliefs.

Because we use beliefs as a richer indicator of subjects' behavior, we can investigate their way of reasoning while avoiding a number of restrictive assumptions that are usually made in studies based on choices. Many behavioral solution concepts induce a probability distribution over choices requiring mixture models to assess subjects' strategic thinking. Stahl and Wilson (1994, 1995) are the first to employ mixture models to infer subjects' approaches to strategic thinking, followed by subsequent work of Costa-Gomes et al. (2001), Costa-Gomes and Crawford (2006), Weizsäcker (2003), Costa-Gomes and Weizsäcker (2008) and others. Mixture models rely on the assumption that subjects who decide according to the same concept also share the exact same parameter values for the given concept and that these parameters are constant across games. Our method avoids these assumptions. We simply check if a belief falls within the area of beliefs predicted by a given concept, where the area is generated by considering all possible parameter values for the concept. If it does, we can infer that a subject's beliefs are in line with the concept. Thus, we can assess whether a subject is consistent with a concept without any assumptions on the parameter

values. We use Selten's measure of predictive success (MPS) (Selten and Krischker, 1983, Selten, 1991) to reveal the concept which explains subjects' behavior most efficiently. Using parametric bootstrap procedures we can test if a given theory performs better than random and if a subject states beliefs as if he reasons consistently according to a concept across games.

We find that for about 80 percent of our 88 subjects the concept with the highest MPS for a given subject explains more beliefs than random beliefs would. The concept with the highest MPS is L1 for 7 subjects, L2 for 21, L3 for 23, UPB for 15, and QRE for 2 subjects. Furthermore, we test if there are subjects who state beliefs as if they reason according to a given concept across all games with some room for mistakes, which we call consistency. We find that 45 percent of our subjects are consistent with a concept, whereof 4 are consistent with L1, 13 with L2, 20 with L3, and 2 with UPB. Overall, the concepts have explanatory power for a majority of subjects and describe half of them with consistency.

To obtain further insights about subjects' reasoning approaches, we analyze the time they take to state their beliefs. For Lk we can form intuitive hypothesis about the relative response times.³ Level-k models assume that the higher a person's level of reasoning (i.e., the higher his k), the more steps of reasoning he makes (Nagel, 1995). Thus, we expect that L1 agents take less time to state their beliefs than L2 agents and L2 agents less than L3 agents. And indeed, this hypothesis can be confirmed for subjects who reason according to Lk. This evidence further supports our finding that these subjects do follow the Lk logic.

The predictive power of beliefs is of course crucial for the relevance of our analysis. Standard subjective expected utility theory (Savage, 1972), which is the basis for most economic and game theoretic models, assumes that agents best-respond to their beliefs. This assumption is relaxed for beliefs as well as choices in all three concepts we discuss. Both QRE and Lk assume logit responses, when forming beliefs as well as when making choice. Regarding choices logit responses are sensitive to the differences in expected payoff induced by the belief. For example, if the difference in expected payoff between two choices is very small then the logit map will assign almost equal probability to the two choices,⁴ whereas the best-response mapping will put full weight on the optimal choice. UPB's utility proportional mapping works similarly to the logit map, only the probabilities assigned to the opponent's choices are linearly proportional to the

³For the other concepts there is no clear hypothesis about the relative response times.

⁴For small parameter values.

differences in utilities. To test if this kind of behavior is present in our sample we ran a logistic regression analysis with best-response as the dependent variable. We find that the difference in expected payoff induced by the stated belief plays an important role. If a belief clearly favors one choice, our subjects are a lot more likely to best-respond. Moreover, we see that more than half of the observations fall within an interval where best-responding might be computationally difficult. Within this interval the chance of best-responding is around 50 percent. However, if the difference in expected payoff becomes sufficiently large the chance of subjects best-responding increases to around 85 percent. We conclude that beliefs are important for choices but might not be deterministic, especially if finding the best choice is computationally difficult. These insights on the relation between beliefs and choices strengthens our trust in a classification on the basis of beliefs.

We proceed by first discussing the concepts under investigation in Section 3.1. In Section 3.2 we elaborate on the experimental design and the specific characteristic of the games. Then we present our results including robustness checks in Section 3.3 and finally discuss their connection to the literature in Section 3.4.

3.1 Concepts

Throughout the paper we will discuss a variety of concepts and games. Therefore, we introduce the notation here independently of the concepts and games. Let I denote a the set two players i and j , C_i denotes player i 's finite choice set and let $C := C_i \times C_j$. Player i 's payoff function is denoted by $U_i : C \rightarrow \mathbb{R}$. Moreover, player i can choose from $|C_i|$ choices. Player i 's belief about player j we denote by $p_i \in \Delta(C_j)$ and by $p_i(c_j)$ the probability i assigns to j 's choice c_j . Player i 's expected payoff function is then given by $u_i(c_i, p_i) = \sum_{c_j \in C_j} p_i(c_j) U_i(c_i, c_j)$ and player i 's average payoff by $u_i^{avg}(p_i) := |C_i|^{-1} \sum_{c_i \in C_i} u_i(c_i, p_i)$. We define $\bar{u}_i := \max_{c \in C} u_i(c)$ and $\underline{u}_i := \min_{c \in C} u_i(c)$. Our analysis is mostly based on the players' beliefs, hence we will make a clear distinction between what we consider as beliefs and what we will consider as choices. By p_i we denote beliefs about the opponent's choices and by q_i a player's stochastic choice function, where $p_i \in \Delta(C_j)$ and $q_i \in \Delta(C_i)$.

The noisy level-k model and the logit QRE share the logit mapping as main ingredient. The mapping depends on a noise parameter $\lambda_i \geq 0$ for every player i ,

which specifies the level of noise that enters the response. As λ_i approaches zero, the logit mapping assigns equal probability to all the opponent's choices. On the other end of the spectrum, as λ_i approaches infinity the logit mapping becomes a best-response function. Formally the logit mapping is defined as follows:

$$p_i(c_j, p_j, \lambda_i) := \frac{\exp \{ \lambda_i u_j(c_j, p_j) \}}{\sum_{c'_j \in C_j} \exp \{ \lambda_i u_j(c'_j, p_j) \}}. \quad (3.1.1)$$

As already discussed, L0 players are assumed to play all actions with equal probability. Therefore, we assume that q_i^{L0} is a stochastic choice function, where player i plays all choices with equal probability. A L1 player believes that his opponent plays all choices with equal probability. His beliefs are denoted by p_i^{L1} , where $p_i^{L1} := q_j^{L0}$. The beliefs of a L2 player can be obtained by $p_i^{L2} := p_i(c_j, p_j^{L1}, \lambda_i)$, the beliefs of a L3 player by $p_i^{L3} := p_i(c_j, p_j^{L2}, \lambda_i)$, and so on. We assume that the noise, i.e., the λ_i , is constant in a player's belief formation. For $k \geq 1$, a Lk player's stochastic choice is given by $q_i^{Lk}(c_i, p_i^{Lk}, \gamma_i) := p_j^{L(k+1)}(c_i, p_i^{Lk}, \gamma_i)$. The parameter $\gamma_i \geq 0$ is independent from λ_i and describes the noise that enters the stochastic choice. Therefore mistakes can occur in the same way when forming beliefs and choices. We define the stochastic choice functions for the other levels accordingly. Note that a level-k player's beliefs are equivalent to the stochastic choice function of a level-(k-1) player.

The logit QRE is a noisy equilibrium in stochastic choices and is obtained as a fixed point of the $q_i \circ q_j$ mapping:

$$q_i^{QRE}(c_i, q_j^{QRE}, \lambda_i) := \frac{\exp \{ \lambda_i u_i(c_i, q_j^{QRE}) \}}{\sum_{c'_i \in C_i} \exp \{ \lambda_i u_i(c'_i, q_j^{QRE}) \}}, \quad (3.1.2)$$

where q_i^{QRE} and q_j^{QRE} constitute a fixed point of the $q_i \circ q_j$ mapping. Hence, player i responds to j 's stochastic choice. This can be interpreted as player i believing that player j will make choices where their probability is given by j 's stochastic choice function. Thus, we can obtain player i 's beliefs via

$$p_i^{QRE}(c_j, p_j^{QRE}, \lambda_i) := q_j^{QRE}(c_j, q_i^{QRE}, \lambda_i),$$

i.e., player i 's belief about j is j 's stochastic choice function.

UPB assumes that differences in probabilities that a player assigns to his op-

ponent's choices is proportional to the expected utilities they generate

$$p_i(c_j, p_j) - p_i(c'_j, p_j) = \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} \left(u_j(c_j, p_j) - u_j(c'_j, p_j) \right)$$

where $c_j, c'_j \in C_j$. The probability that player i assigns to player j 's choice c_j can be obtained by

$$p_i(c_j, p_j) = \frac{1}{|C_j|} + \frac{\lambda_j}{\bar{u}_j - \underline{u}_j} \left(u_j(c_j, p_j) - u_j^{avg}(p_j) \right), \quad (3.1.3)$$

where $\lambda_j \in [0, \lambda_j^{max}]$ is player i 's sensitivity to player j 's differences in utilities. The maximal sensitivity λ_j^{max} is chosen such that equation (3.1.3) always yields well-defined probabilities.⁵ Bach and Perea formalize this concept using an epistemic model and common belief in UPB to characterize the beliefs a player can hold. Furthermore, Bach and Perea show that the beliefs players' can hold under UPB are unique in the two player case. Hence, UPB gives unique beliefs for the class of games we consider.

In two player 2×2 games UPB is technically a QRE. However, in games with more choices or players this is not true. Therefore, different predictions in beliefs of the logit QRE and UPB in our 2×2 games are solely due to the different mappings they employ. Logit QRE uses the logit map in (3.1.1) whereas UPB uses (3.1.3), which is linear in probabilities. Figure 3.1.1 shows how depending on the value of λ_i the impact of differences in utilities varies.⁶ For small values of λ_i small differences in utility have a stronger impact on the differences in probabilities whereas for small values of λ_i any difference in utilities has a weaker effect than in the linear case.

UPB is also very close to L2 when setting the sensitivity parameter for a player's own payoffs to zero. To avoid that L2 and UPB coincide, we assume throughout the paper that $\lambda_i > 0$ to isolate the important characteristics of UPB.

⁵Since this assumption can limit the range of UPB's predictions and it is a mathematical rather than a behavioral assumption, we used non-linear optimization to find λ^{max} such that every step in the corresponding reasoning process (Nauerz, 2015) is well-defined. Note that Chapter 2 is based on Nauerz (2015).

⁶The graph is a plot of equation (3.A.1) and can be found in Appendix 3.A.

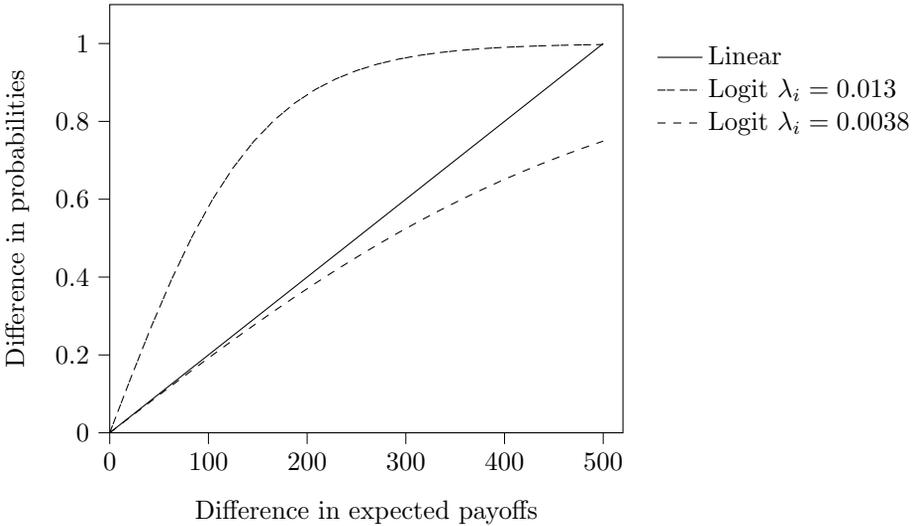


Figure 3.1.1: Differences in probabilities as function of differences in expected payoffs

3.2 Experimental Method

In the main part of our experiment, subjects stated beliefs and made choices in 28 one-shot bi-matrix games without feedback in between games.⁷ The games were always presented from the perspective of the *Row Player*. Each game assigns a payoff to each of the two players depending on the combination of their choices (see Figure 2.5.1 for an example of a game). For each game, subjects first stated their belief about the *Column Player's* choices and then made their own choice.⁸

3.2.1 Procedures

The experiment took place in the *BEElab* at Maastricht University in February 2015. We recruited 88 subjects through the ORSEE subject pool (Greiner, 2004). The experiment was computer based and programmed with ztree (Fischbacher, 2007). We expressed subjects' payoffs in Experimental Currency Units (ECU), where 12 ECU equaled to 1 EUR. Each of the three sessions lasted about two

⁷In the 2x3 games subjects only had to indicate their belief in case their opponent had 2 choices.

⁸One might be inclined to argue that the belief elicitation has effects on the subjects' choice behavior. However, Costa-Gomes and Weizsäcker (2008) find that belief elicitation, even right before the choice, does not induce significantly different beliefs than eliciting the choice first.

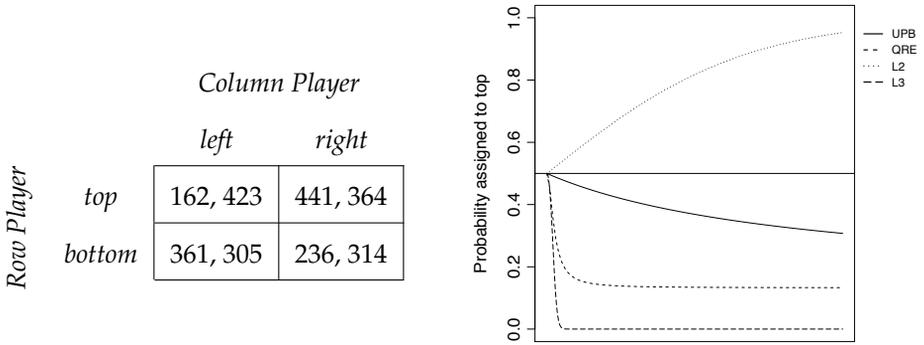
hours. It was common knowledge that subjects remain anonymous towards other subjects and the experimenter throughout the whole experiment. After reading the instructions, subjects answered a series of control questions to ensure that they understood the instructions correctly. Three subjects answered at least one control question incorrectly. After explaining the corresponding part of the instructions again, they answered the respective control questions anew. All three subjects eventually answered the questions correctly. Subjects' questions related to the experiment were answered privately in the cabins. At the end of the experiment, subjects filled in a questionnaire. Subjects received a fixed payoff of EUR 2.50 and earned an average of EUR 29.50 in the belief/choice task.

9

3.2.2 Games

Our games are 2×2 and 2×3 strictly competitive games designed in such a way that putting more weight on one of two choices can only be explained by one or two concepts. This feature provides enough variance in the MPS between concepts to distinguish subjects' approach to strategic thinking. When identifying subjects' approach, the type of games is also an important issue. The games' nature might trigger certain kinds of thinking. For example there is a stream in the literature that investigates strategic thinking in 2 person guessing games (e.g. Costa-Gomes and Crawford (2006), Fragiadakis et al. (2013), or Georganas et al. (2014)). Guessing games, might invoke the search for a starting point to evaluate actions—a feature that 2×2 and 2×3 strictly competitive games are unlikely to possess. Even if strictly competitive games have the same effect on players' reasoning, there exists no obvious starting point. Thus, there exists no clear level-k structure as for example in the beauty contest. There also exists no equilibrium in pure strategies that could serve as an anchor or dominated choices, which could be eliminated. Hence, we believe that this type of games offers as little structure as possible and thus gives a more reliable indication of the subjects' approach to strategic thinking. Moreover, their simple structure makes assessing beliefs a simple task as it boils down to evaluating weights between two options.

⁹The average payoff is for stating beliefs and making choices in 32 games. After making the 28 decisions in the main part of the experiment, subjects played 4 of the 28 games again. Half of subjects received an additional EUR 0.50 per game if they remembered a 7-digit binary number correctly while playing those games. These additional games were part of a treatment which is not discussed here. In this paper, we focus on stated beliefs and choices in the first 28 games.



(a) Game #2 (game type L2). Payoffs in the game are in ECU. The game assigns a payoff to the *Row Player* and the *Column Player* depending on the combination of their choices. For example, if the *Row Player* chooses *bottom* and the *Column Player* chooses *left*, then the *Row Player's* payoff is 361 ECU and the *Column Player's* payoff is 305 ECU.

(b) Probability assigned to *top* as function of the parameter value under L2, L3, UPB, and QRE. The parameter for L2, L3, and QRE is the noise parameter, and for UPB it is a player's sensitivity to her opponent's differences in utilities. The parameter values for the different concepts are scaled differently for the sake of presentation.

Figure 3.2.1: Example game

Take the game in Figure 2.5.1 as an example.¹⁰ The concepts' predictions as function of their respective parameter value are given in Figure 3.2.1b. Consider the *Column Player's* belief about the *Row Player*. In this example assigning a higher probability to the *Row Player's* choice *top* can only be explained by L2 for all possible parameter values for all concepts. The other concepts predict that the *Column Player* would deem the *Row Player's* choice *bottom* more likely regardless of the respective parameter value. Thus, if a subject states a belief that puts more weight on the *Column Player's* choice *top* her stated belief is in line with the concept L2. For a complete overview of all 28 games, see Appendix 3.C.

Table 3.1 shows the different types of games we use to differentiate the concepts. Game type L2 and game type UPB enable us to isolate L2 and UPB, respectively. In game type L3&QRE, both concepts make the same prediction about beliefs, but QRE-beliefs remain very close to 0.5 for any parameter value, while

¹⁰The reader might be concerned that 3-digit numbers influence the subjects' behavior. However, Rey-Biel (2009) found no significant influence on subjects' behavior in 3×3 normal-form games.

Table 3.1: Overview of Game Types^a

Game type	Isolated	Size	No. of Games
L2	L2	2×2	4
L3&QRE	L3, QRE	2×2	4
UPB	UPB	2×3	4
UPBvQRE	-	2×2	2

^a The 4 game types each isolate different concepts. The forth column indicates for each game type how many games a subject played in which the given concept(s) could be isolated for a subject.

L3-beliefs quickly move away from 0.5 beliefs as the parameter value increases. This enables us to distinguish between QRE and L3 players. Since UPB is also a QRE in 2×2 games¹¹ we use *UPBvQRE* games to investigate which specification prevails. Here UPB puts more weight on a different choice than QRE for all parameter values. For a detailed discussion of why the concepts select certain choices in each game type consult the Appendix 3.B. During the course of the experiment every player should face four *L2* games, four *L3&QRE* games, four *UPB* games and two *UPBvQRE* games. A single game allows us to make inferences about one of the two players. Thus, every subject needs to play a total of 28 games. Note that we do not need a separate game to identify *L1* players as only exact 0.5 beliefs indicate that a subject reasons in line with *L1*.

3.2.3 Belief Elicitation and Incentives

To elicit the probability subjects assign to their opponent's choices, subjects were asked the following question:

"How likely do you regard the event that the OTHER PLAYER chooses RIGHT?"

Subjects indicated their answer on a scale from 0 (totally unlikely) to 100 (totally likely) in multiples of 1 by means of a slider on the computer screen. Figure 3.2.2 shows a screenshot of the belief elicitation stage in ztree. Depending on which likelihood subjects chose to position the slider, their payoffs in case their opponent chooses *left* and in case he chooses *right* were shown. They could reposition

¹¹See Section 3.1.

the slider as many times as they wanted before finalizing their choice. Additionally, subjects received a payoff table in the instructions displaying their payoffs depending on their opponent's choices for each possible probability statement (for the full instructions, see Appendix 3.E). We asked subjects' beliefs only for those games in which their opponent had two choices, i.e., in 24 games.

In order to obtain subjects' true probabilistic beliefs, the payoffs were constructed using the L-adjusted Quadratic Scoring Rule (LQSR) (Offerman and Palley, 2015). The LQSR is based on a prospect theory model of Palley (2013) and accounts both for subjects' risk and loss aversion. In this way, the LQSR is designed to eliminate excessive reporting of beliefs of 1/2, which is commonly observed when using the Quadratic Scoring Rule (QSR) (e.g. Offerman et al., 2009). In the case of the QSR it can be optimal for a loss averse agent to state a belief of 1/2 even if his true beliefs are close to 1/2. In the LQSR the adjustment factor calibrates the scoring rule such that compared to the baseline probability of 1/2 an ex post belief in the wrong direction is punished less severely than an ex post belief in the correct direction is rewarded. We use the adjustment factor $L^* = 3.7$ for the general population.¹² Offerman and Palley (2015) show experimentally that this adjustment performs well in obtaining subjects' true beliefs.¹³ Box 2 in the instructions of Appendix 3.E shows subjects' payoffs for the range of possible stated beliefs for both choices of their opponent. For example, if a subject states that he believes her opponent will choose *right* with a probability of 52 percent, he will receive a payoff of 256.94 ECU if his opponent indeed chooses *right* and 250.61 ECU if his opponent chooses *left*. Hence, compared to the baseline probability of 50 percent with a risk free payoff of 252 ECU, stating a belief in the correct direction (i.e., ex ante deeming the choice that the opponent chooses in the end more likely) means gaining an additional 4.94 ECU. On the other hand, stating a belief in the wrong direction (i.e., ex ante deeming the choice that the opponent chooses in the end less likely) means losing only 1.39 ECU. In this way, even for a loss averse subject it can be optimal to state his true belief of 52 percent.

At the beginning of the experiment, subjects were randomly assigned the role of either the *Row Player* or the *Column Player*. Each subject was paid either

¹²We do not apply individual adjustment factors, since this would entail an additional choice task before playing the games. This task could potentially bias subjects' choices and belief statements.

¹³They estimate the goodness of subjects' stated beliefs by eliciting them using the LQSR with different adjustment factors. $L^* = 3.7$ performs significantly better than the QSR without adjustment and also (insignificantly) better than individual adjustment factors obtained in a separate choice task.

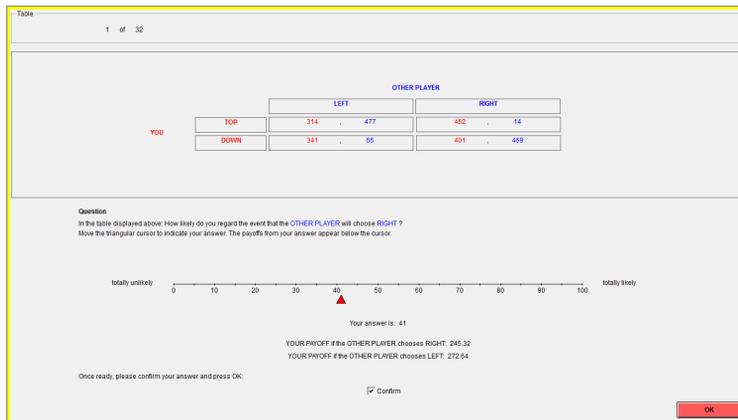


Figure 3.2.2: Screenshot of belief elicitation stage in ztree

according to a belief or a choice in one of the games. Both selections were made at random to mitigate hedging behavior (Blanco et al., 2010).¹⁴ Subjects were then randomly matched with a subject of the opposite role. If a belief was selected, a subject's payment was determined by her stated belief using the LQSR and the choice of her assigned opponent in the respective game. If a choice was randomly drawn, her payment was determined by the outcome of the respective game.

To avoid confounding effects of social considerations, we instructed subjects to maximize their own payoff. We also informed them that all other subjects were given this instruction (Dreber et al., 2013).¹⁵ As far as risk aversion is concerned, remember that the LQSR ensures that risk attitudes do not influence stated beliefs. Moreover, note that constant relative risk aversion cannot influence players' choices. If we incorporate constant relative risk aversion (as suggested by Goeree et al. (2003)) in the logit response function then we simply obtain a monotone transformation of the function, which does not alter the choice.

¹⁴The game was selected from the total of 32 games subjects played.

¹⁵As Dreber et al. (2013) we find that the manipulation indeed induced subjects to maximize their own payoff. At the end of the experiment we asked our subjects if their concern for their opponent's payoff influenced their choice or if they thought that their opponent's choice might be influenced by a concern about their payoff. For both questions, about 70 percent of subjects indicated that they thought this was the case "in no tables" or "in hardly any tables". Furthermore, we find that classified subjects are less influenced by other regarding preferences in both beliefs and choices on a 10 percent significance level.

3.3 Results

This section presents the outcomes of our experiment and the analysis thereof. We first present some descriptive evidence about subjects' stated beliefs, in Subsection 4.1. Then, in Subsection 4.2, we study whether the concept with the highest MPS performs better than random choice. In Subsection 4.3, we investigate if there are subjects who consistently reason according to a single concept across games. To provide secondary support for our classification, Subsection 4.4 presents an analysis of subjects' response times. Finally, in Subsection 4.5 we discuss the connection between beliefs and choices in our experiment.

3.3.1 Descriptive Evidence

First of all, we present an overview of the data on subjects' beliefs. Figure 3.3.1 shows the cumulative distribution of the number of games in which subjects stated beliefs in line with a given concept. In other words, it indicates the fraction of subjects who hold beliefs in line with L1, L2, L3, QRE or UPB in at most n games. A belief is considered to be in line with a concept in a given game if it is between the minimum and the maximum belief that can be explained by the concept for any parameter value. Being in line with L1 is defined as having a belief exactly equal to 0.5. As can be seen in Figure 3.3.1, subjects' beliefs are in general more often in line with L2 and L3 than with QRE and UPB, and least often with L1. Four subjects were never in line with UPB, seven never with QRE. About 40 percent of the subjects never stated a belief of 0.5 and were therefore never in line with L1. Every subject stated a belief in line with L2 at least 12 times and with L3 at least 11 times.

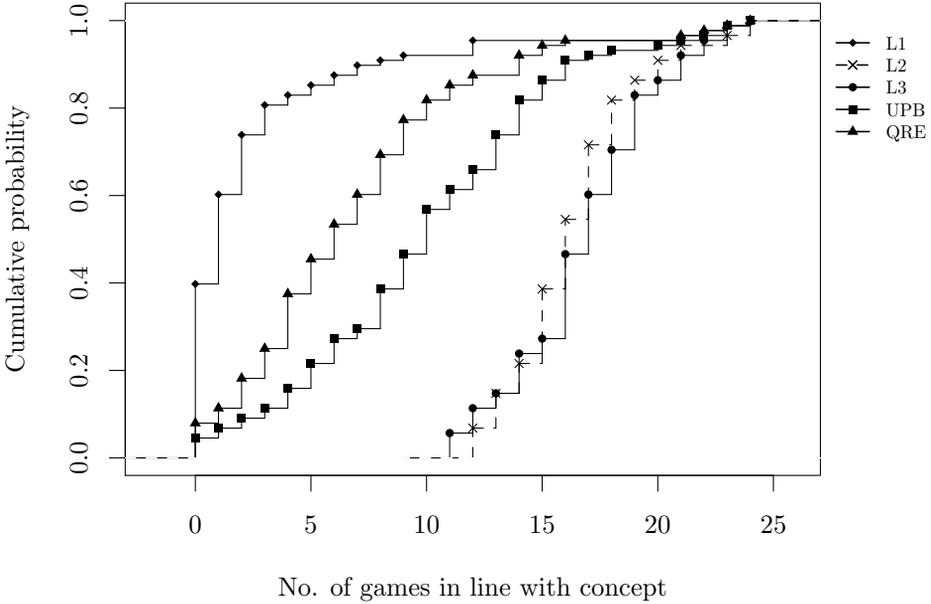


Figure 3.3.1: Cumulative distribution function of concepts

The results presented above are based on the beliefs stated by subjects in 24 games. Among those games, we have for each player four "isolation games" for each concept (described in Section 3.2). In these games certain beliefs can only be explained by a single concept. Hence, if a subject's strategic thinking is similar to one of the concepts he should state beliefs in line with his approach in most of the isolation games and also over all games. Figure 3.3.2 illustrates that individuals with at least three beliefs in line with a given concept in the respective isolation games also seem to follow this concept clearly more often in general. A series of one-sided t-tests confirms that for each concept, those individuals who follow the concept at least three times in the isolation games also follow the concept significantly more often than other subjects in all games (at a 10-percent significance level). This suggests that the performance of a concept in the isolation games is strongly correlated with its overall performance. However, this view does not take into account the efficiency of the concept at hand.

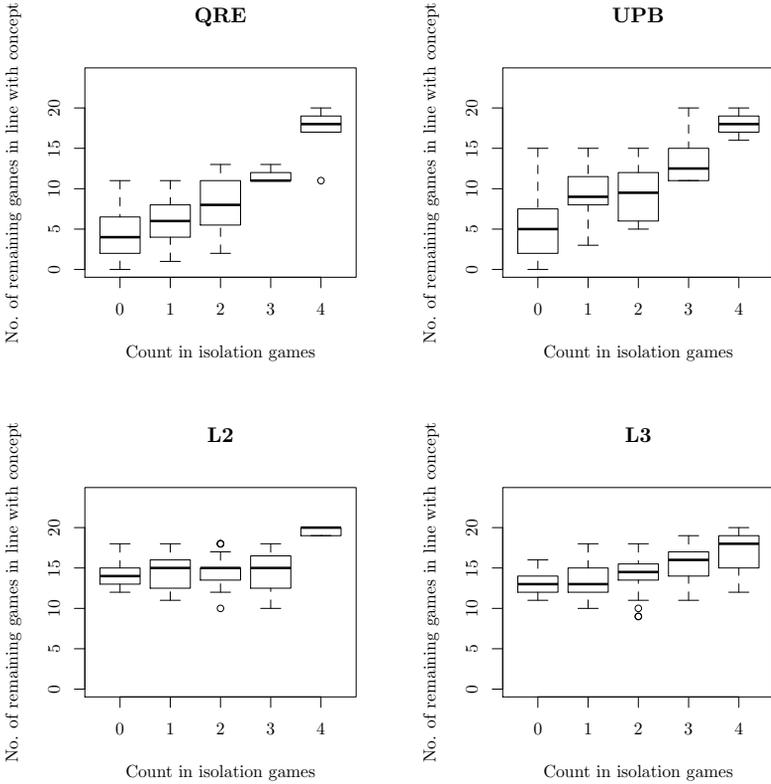


Figure 3.3.2: Relation between isolation games and overall consistency with a concept. For each concept, QRE, UPB, L2, and L3 the x-axis indicates the subjects' count in the isolation games and the y-axis the number of remaining games that are in line with the given concept. Subjects that stated beliefs agreeing with QRE in all four isolation games, for example, stated on average beliefs agreeing with QRE in 18 of the 20 remaining games.

Table 3.2: MPS-based Classification of Subjects^a

Concept	No. of Subjects
L1	7
L2	21
L3	23
UPB	15
QRE	2
Unclassified	20
Total	88

^a Subjects are classified as reasoning according to a given concept if the respective MPS is higher than the MPS scores from the remaining concepts and if the MPS is significantly better than one resulting from random choice. 68 subjects can be classified, 20 remain unclassified and are labeled as "Unclassified."

3.3.2 Performance Against Random Choice

The distributions presented above and the relation of the performance in the isolation game with the overall performance give a first idea of how well the different concepts perform in explaining subjects' stated beliefs. However, comparing raw data might not be the best method, because concepts such as QRE or UPB make predictions that are more restrictive than L2 or L3. For any parameter value, the beliefs predicted by L2 or L3 cover about half of the range of all possible beliefs (i.e., the beliefs of a subject can take any value between 0.5 and either 0 or 1). QRE and UPB each predict a more limited range of beliefs, that typically covers about a quarter of the range of all possible beliefs (for an example, see 3.2.1b). In order to assess the performance of the concepts, we need to take these differences into account. To do so, we use the measure of predictive success proposed by Selten and Krischker (1983) and axiomatized by Selten (1991). It compares the hit rate of a theory (i.e., the relative frequency of correct predictions) with the area (i.e., the share of all possible outcomes that is predicted by the theory). Our design is particularly well suited to use this measure. First, our games enable us to clearly separate between concepts on the basis of beliefs, so that the hit rate has sufficient variation across the different concepts. Second, the range of beliefs that a concept can predict in a game regardless of the parameter value exactly fits the idea of the area. Thus, by using Selten's MPS, we can select the theory which explains subjects' behavior most efficiently.

To calculate the MPS for a given concept and subject $s \in \{1, \dots, 88\}$ in game

$h \in \{1, \dots, 24\}$ and role $i \in \{Row, Column\}$, we simply check whether or not the observed belief is within the area predicted by the concept. This is encoded in a binary variable called r_{sh} . The size of the area a_{hi} is the area covered by beliefs that can be explained by a given concept in that game. It ranges from 0, if no belief can be explained by the concept, to 1, in case the concept can explain all beliefs. The score is calculated as:

$$m(s, h, i) := r_{sh} - a_{hi}.$$

We construct an MPS per subject for each of the concepts by averaging scores over the 24 games in which the subject had to state a belief.

By comparing the MPS we can find the best-performing concept for every subject, i.e., the concept that has the most predictive power after correcting for the area of predictions. Since a positive MPS can also be generated by random choice, we perform a hypothesis test to exclude this possibility. Because we have a small sample we use a parametric bootstrap to construct the distributions of the MPS that would result from random choice. To simulate one observation, we first draw 24 random beliefs from a uniform distribution over $[0, 1]$ and then calculate the resulting MPS for each concept and game. Finally, we aggregate over the games. This step is then repeated 1000 times to generate a random distribution for each concept. We use this procedure to test whether the MPS of the best-performing concept is significantly higher than the MPS produced by random beliefs. The one-sided hypothesis tests are performed at a 5-percent confidence level. Table 3.2 shows for how many subjects each concept is best-performing and at the same time significantly better than random. This is true for 7 subjects in case of L1, for 21 subjects in case of L2, for 23 subjects in case of L3, for 15 subjects in case of UPB, and for 2 subjects in case of QRE. For 20 subjects, the best-performing concept does not score significantly better than random.

3.3.3 Performance Against Perfect Consistency

A further indicator that a subject is reasoning in line with a concept is whether he does so consistently across games. We define consistency as following a concept in all games with some room for mistakes. Thus we need to determine the distribution of the MPS for a subject who is consistent with a concept. This allows us to test whether the observed MPS for each subject is significantly different from this benchmark distribution of hits. In an ideal world, subjects have a

fixed approach of strategic thinking from the beginning of the experiment and are never distracted. Then, consistency with a concept would mean that all stated beliefs can be explained by the concept. The subjects would never make mistakes, and the hit rate would be maximal. However, this does not appear to be a realistic benchmark in the case of an experiment with subjects who are inexperienced with the task. We have to take at least two factors into account which might have caused subjects to make mistakes. First, it might be that subjects had not fixed their approach from the very beginning of the experiment. Second, the experiment lasted for about two hours. This makes it possible that subjects were distracted in some games and provided an answer which is different from what they would have answered had they paid full attention. Therefore, we model the number of mistakes made by a subject using a Poisson distribution with a mean-variance parameter of three. The choice reflects our best assessment and is also comparable with the number of mistakes estimated in Costa-Gomes et al. (2001). Robustness checks show that our results are mostly stable even for lower values of the mean-variance parameter (see Table 3.6 in Appendix 3.D).¹⁶

To obtain the benchmark distribution of hits, we again use parametric bootstrapping. We draw a number of mistakes from the Poisson distribution. We then attribute these mistakes to games at random, and eventually calculate the resulting MPS. We repeat this process 1000 times in order to obtain a benchmark distribution of the MPS for each concept. For each subject, we can then test whether the MPS of the best-performing concept is significantly lower than the one of a subject who is consistently following the concept. Table 3.3 presents the number of subjects for which we cannot reject the null hypothesis of being consistent with a given concept. We find that 4 subjects are consistent with L1, 13 with L2, 20 with L3, and 2 with UPB. A bit more than half of our sample does not appear to be consistent with any concept of strategic thinking across games.

3.3.4 Response Times

Fortunately, computerized experiments give us the opportunity to also take secondary information about subjects' behavior into account. We check whether our classifications are meaningful by comparing how long classified individuals needed to state their beliefs. Lk models assume that the higher a person's k , the

¹⁶For a mean-variance parameter of 2 only the number of subjects classified as L2 varies substantially, suggesting that subjects who reason according to L2 make more mistakes than subjects who reason according to L1, L3, and UPB.

Table 3.3: Subjects Classified as Consistent^a

Concept	No. of Subjects
L1	4
L2	13
L3	20
UPB	2
QRE	0
Unclassified	49
Total	88

^a Subjects are classified as consistently following a concept if the respective MPS is higher than the MPS scores of the remaining concepts and if we cannot reject the hypothesis that a subject states beliefs as he would if he truly followed the concept. 39 subjects can be classified, and 49 remain unclassified and are labeled as "Unclassified."

more steps of reasoning he makes. Each additional reasoning step increases the time it takes to form the final belief. Therefore, for L_k players, we hypothesize that the higher the *k*, the longer the time needed to state their beliefs. L1 players assume their opponents randomize, hence they state the uniform belief. L2 players believe that their opponent believes that they play uniformly, a more complicated process than assuming random play. And finally, L3 players undertake a further step of reasoning. They believe that their opponent believes that they believe that he will play uniformly. Since there exists no clear hypothesis about the relative response times of L_k, UPB, and QRE players, we focus on the analysis of the differences between L1, L2, and L3 players.

We recorded individual response times in the belief elicitation stage of each game.¹⁷ In Table 3.4, we regressed the time it took subjects to state their beliefs on the concept with which they are consistent, while controlling for the type of game they were considering and the period. L2 is used as the reference category. When subjects are classified according to the highest MPS, subjects classified as L1 take on average 12.2 seconds less than subjects classified as L2. Subjects classified as L3, on the other hand, need on average 8 seconds longer than subjects classified as L2. The same pattern is also present when we classify subjects according to consistency. To summarize, these findings are statistically significant and in line with the hypothesis stated above. Hence, our results provide sug-

¹⁷Response times were recorded by ztree as the time which passes from the moment that a subject sees the screen for stating his belief in a given game until he presses the "OK" button.

Table 3.4: Response Times OLS Estimates^a

	Response Time Belief	
	(1) Highest MPS	(2) Consistent
Class L1	-12.200***	-12.446***
Class L3	8.057***	12.731***
Class Unclassified	-0.198	6.974***
Class QRE	21.778***	
Class UPB	1.434	-0.359
Game L2	-3.442**	-3.442**
Game UPBvQRE	-5.018***	-5.019***
Game UPB	-10.839***	-10.846***
Period	-1.270***	-1.271***
Constant	56.151***	51.795***
Observations	2,112	2,112
R ²	0.185	0.180
Adjusted R ²	0.182	0.177
Residual Std. Error	24.578 (df = 2102)	24.656 (df = 2103)
F Statistic	53.188*** (df = 9; 2102)	57.675*** (df = 8; 2103)

^a The time subjects' took to state their beliefs regressed on classifications due to (1) highest MPS, (2) consistency, and (3) isolation games as well as the type of game and the period. For example when we consider the classification due to consistency, we see that L1 players are on average 12.6 seconds faster compared to subjects classified as L2 when stating their beliefs. L3 players, on the other hand, take 12.6 seconds longer than subjects classified as L2. Significance levels are reported as usual: *** denotes the 1 percent level, ** denotes the 5 percent level, and * denotes the 10 percent level.

gestive evidence for the general idea of the reasoning processes behind Lk and therefore that our classification using the MPS is meaningful.

3.3.5 Best Responses

Since we base our classification on subjects' beliefs, it is important to discuss the connection between beliefs and choices. Clearly, our classification would be meaningless if beliefs had no influence on choices. We commonly assume that an individual's belief induces an expected value over his choices and that an individual chooses such that he maximizes this expected value. Therefore, if beliefs reflect the true assessment of a situation individuals lose payoffs by not best-responding to their beliefs. Potential losses due to best-response mistakes are equal to the distance between the expected values associated with each of

the choices, given a player's beliefs in the game. Figure 3.3.3 presents the cumulative distribution function of losses due to best-response mistakes made by individuals in our experiment, along with the cumulative distribution due to random choices as a benchmark. We see that subjects in our experiment do substantially better than random choice. On average they best-respond in more than 65 percent of the cases. Moreover, we see that they make a lot of mistakes when the loss is small (≤ 100 ECU) and few when the loss is high (> 100 ECU). Over 90 percent of the observations contain choices that resulted in a loss of less than 100 ECU. Thus, we conjecture that subjects noisily best-respond to their beliefs, so that they best-respond most of the time, sometimes make mistakes, but tend to avoid costly ones.

The idea that individuals tend to avoid costly mistakes is further supported by studying the determinants of the probability of best-responding. Table 3.5 presents the estimation results of a logit model with best-response as the dependent variable and different categories of potential losses due to best-response mistakes as the main regressors. We estimate the model (1) for all subjects and (2) for subjects who are classified due to their consistency with a concept. The results for all subjects indicate that, for instance, moving from a potential loss of less than 50 ECU (the reference category) to a potential loss of 100 to 200 ECU multiplies the odds of best-responding by about 2. The average predicted probability of best-responding jumps up from about 50 percent when the potential loss is smaller than 50 ECU to about 70 percent when the potential loss is between 100 ECU and 200 ECU, and 80 percent when the potential loss is between 200 ECU and 300 ECU. The same pattern holds for the consistent subjects.¹⁸

To summarize, in our experiment, subjects are likely to make best-response mistakes relatively often, because these mistakes often do not have important consequences in terms of payoffs. Therefore, using subjects' choices to classify them in terms of their approaches to strategic thinking is unlikely to yield reliable results. In this case, beliefs can provide a better identification of an individual's strategic thinking.

¹⁸Only the increased likelihood of best-responding for potential loss in ECU between 300 and 400 as well as between 400 and 500 is not significant. This, however, is very likely due to the lack of observations in these categories. There are 39 observations between 300 and 400 and only 14 between 400 and 500 compared to 103 and 41 observations for all subjects.

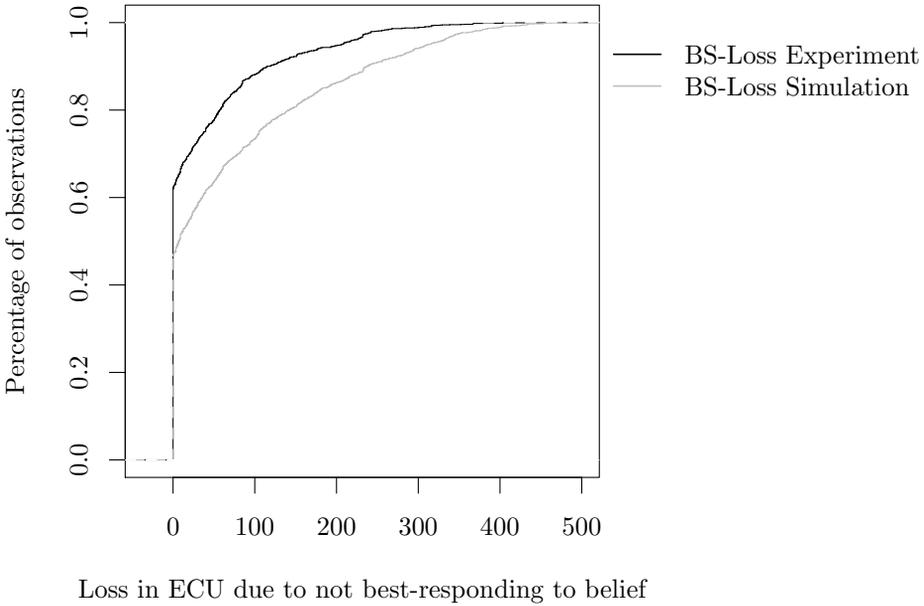


Figure 3.3.3: Cumulative distribution function of losses due to best-response mistakes. Losses in ECU in case a subject doesn't best-respond to his stated belief. A subjects best-responds to his belief if she chooses the option which yields her the highest expected value given her stated belief. $N=2,112$ (24 beliefs for each of the 88 subjects).

3.4 Discussion

In this section, we compare our results to those of other studies that also investigate subjects' strategic thinking in games on an individual basis, all of which use choices to identify subjects' approaches. Both Stahl and Wilson (1994, 1995) as well as Costa-Gomes et al. (2001) use mixture models to classify subjects based on their choices in normal-form games.¹⁹ Rey-Biel (2009) classifies subjects according to concepts if they coincide the most with the respective concept in a

¹⁹Stahl and Wilson (1994) (Costa-Gomes et al. (2001)) classify subjects as following a concept if there is a posterior probability higher than 76 (90) percent that a player followed a certain concept. This is the case for 35 of 40 (38 of 45) subjects.

Table 3.5: Probability of Best-Responding^a

	Probability Best Reponse	
	(1) All	(2) Consistent
Game L2	0.496***	0.756***
Game UPBvQRE	0.207	0.240
Game UPB	-0.815***	-0.678***
EV Distance [50, 100)	0.044	0.156
EV Distance [100, 200)	0.682***	1.005***
EV Distance [200, 300)	1.187***	1.502***
EV Distance [300, 400)	0.913***	0.572
EV Distance [400, 500]	1.493***	1.085
Period	0.002	-0.008
Constant	0.032	0.035
Observations	2,112	936
Log Likelihood	-1,321.852	-571.528
Akaike Inf. Crit.	2,663.704	1,163.056

^a The table presents a logit regression with the probability of best-responding as dependent and the type of game, the distance in expected value induced by the subject's belief, and the period as independent variables. The probability of best-responding for all subjects, for example, is increased by a factor of 2 when the expected value distance induced by the subject's belief is between 100 and 200 instead of between 0 and 50. Significance levels are reported as usual: *** denotes the 1 percent level, ** denotes the 5 percent level, and * denotes the 10 percent level.

series of normal-form games.²⁰ All studies classify more than twice as many subjects as L2 than as L1 players. Furthermore, Rey-Biel (2009) finds about the same amount of L3 as L1 players and about the same amount of dominance types as L2 players.²¹ More than 50 percent of his classified subjects appear to be equilibrium players. Stahl and Wilson (1994) find about the same amount of Nash Equilibrium as L1 players and Costa-Gomes et al. (2001) find some dominance and equilibrium players. Studies using guessing games (Costa-Gomes and Crawford (2006) and Georganas et al. (2014)) generally find more L1 than L2

²⁰Rey-Biel (2009) classifies 102 out of 160 subjects. Results are reported only from the constant sum games, as they are more comparable to our games.

²¹Dominance types eliminate dominated choices and then best-respond to a uniform belief over the remaining choices. The literature distinguishes these types depending on how many rounds of elimination they perform.

players.²² They also classify very few subjects as L3 players or as following the dominance concept. Georganas et al. (2014) finds as many subjects following the equilibrium concept as L1 players.

Our results are comparable in the sense that we find more L2 than L1 players, however, we find even more L3 players in both classifications. The fact that we find more L3 types than previous studies can be due to a variety of reasons. The studies discussed above consider games with dominated choices, which is why L3 and Nash equilibrium will often coincide. Nash equilibrium plays a different role in those games since it is also the outcome of iterative elimination of dominated strategies. In our study only the indifference aspect of equilibrium comes into play, i.e., equilibrium in mixed strategies. Thus, some of our L3 types might have been identified as following the equilibrium concept in previous studies. Furthermore, the dominance concept could possibly coincide with L3. Hence, it is not very surprising that we find more L3 players than previous studies.

As mentioned before, compared to other studies a different aspect of equilibrium is tested here. The fact that we find very few QRE players confirms that the indifference component of equilibrium is very unlikely to predict behavior in one-shot games. Even though QRE is a good tool to predict sample frequencies (e.g. McKelvey and Palfrey (1995), Goeree and Holt (2001), or Goeree et al. (2005)) it is cognitively very demanding, since players do not only need to respond to nondegenerate probability distributions but also find a fixed point for the noisy responses. That NE and QRE are not well-suited to predict individual behavior in one-shot settings is well known (Chong et al., 2005, Crawford and Iriberry, 2007) and therefore no surprise. Since we are evaluating beliefs, we add to these findings by concluding that very few players have equilibrium beliefs.

Moreover, the issue of best responses to stated beliefs has been discussed extensively in the literature. For example, Weizsäcker (2003) and Costa-Gomes and Weizsäcker (2008) suggest players' stated beliefs exhibit deeper strategic thinking than their actions. To draw this conclusion Weizsäcker (2003) estimates a QRE with two separate parameters for the player under analysis and his opponent. He finds that the estimated parameter values for the opponent is lower than the parameter value for the player himself. Costa-Gomes and Weizsäcker (2008) estimate players' beliefs from both the belief task and the choice task. They re-

²²Using the players' choices classify 58 of their 88 subjects. Georganas et al. (2014) classify their subjects if they agree with one of the concepts in at least 8 of their 20 games. According to this rule 45 of the 150 subjects can be classified.

ject the hypothesis that both are generated by the same belief. These findings along with other studies that find imperfect responses (e.g. Rey-Biel, 2009) suggest that players might not be best-responding to their beliefs and therefore calls the validity of beliefs as indicator of behavior into question. Our findings on the other hand show that players noisily best-respond and that in more than half of the cases it would be computationally too costly for subjects to best-respond. This means that about half of the *choice* data in our experiment is pure noise and fitting a statistical model to it is unlikely to be informative. Therefore, we conclude that beliefs are the better indicator of subjects' reasoning for our experimental setup.

3.5 Concluding Remarks

The way agents approach one-shot strategic interactions is an important and non-trivial problem in economics. We elicit subjects' beliefs in order to infer their approach to strategic thinking in these settings. Beliefs are preferred to choices as they provide more fine-grained information about subjects' assessment of the situation. We take two different approaches to identify subjects' reasoning without the need to make restrictive assumptions about homogeneity of parameters across subjects and/or games. Both approaches lead to similar conclusions. The majority of subjects reasons according to the noisy level-k model, where every step of the reasoning procedure is a noisy best-response to the previous. Noisy best-responses reflect subjects' computational limitations in the sense that they cannot effortlessly perform perfectly accurate computations, as well as other sources of mistakes, such as distraction. This holds for both forming beliefs and making choices as revealed in our analysis, further supporting logit responses when modeling both beliefs and choice. Yet, there appears to be a significant amount of heterogeneity between subjects.

Our method identifies subjects whose beliefs can be described by a single concept within a fixed class of games. A natural question is whether subjects' behavior is also stable across different classes of games. Georganas et al. (2014) suggest that the behavior across classes of games is not stable. Even a weak assumption as perseverance of sophistication does not seem to hold. Our design might be helpful to further investigate this issue since it employs weaker assumptions and therefore might uncover insights that were previously masked.

3.A Appendix - Logit-QRE as Differences in Utilities

To make an exact comparison between the logit QRE and the UPB map, we express the logit QRE map in terms of differences in utilities. First, we take the ratio between two choices $c_j, c'_j \in C_j$,

$$\begin{aligned} \frac{q_i(c_i, p_i, \lambda)}{q_i(c'_i, p_i, \lambda)} &= \frac{\exp(\lambda u_i(c_i, p_i))}{\exp(\lambda u_i(c'_i, p_i))} \\ &= \exp(\lambda(u_i(c_i, p_i) - u_i(c'_i, p_i))) \end{aligned}$$

Now note, that if there are only two choices for player j , then the difference between the probabilities assigned to these choices is uniquely determined by their ratio. Take $\alpha \in [0, 1]$ as any probability assigned to choice c_j and $\beta := \exp(\lambda(u_i(c_i, p_i) - u_i(c'_i, p_i)))$. Then we can express (3.1.2) as $\alpha/(1 - \alpha) = \beta$ and thereby $\alpha = \beta/(1 + \beta)$. The difference between the two can then be expressed as

$$q_i(c_i, p_i, \lambda) - q_i(c'_i, p_i, \lambda) = \frac{2 \cdot \exp(\lambda(u_i(c_i, p_i) - u_i(c'_i, p_i)))}{1 + \exp(\lambda(u_i(c_i, p_i) - u_i(c'_i, p_i)))} - 1. \quad (3.A.1)$$

3.B Appendix - Description of Games

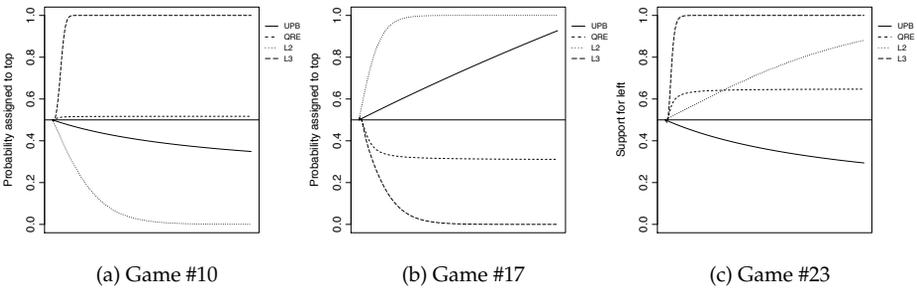


Figure 3.B.1: Concepts' predictions as function of the parameter value

In the following we give four examples that explain how the games separate the concepts predictions in beliefs. All statements are independent of the parameter values of the concepts. Only when we use a converging argument for QRE we assume low parameter values such that the logit mapping is a contract mapping.

We describe how the *Column Player* forms his belief about the *Row Player*.

L2 Games

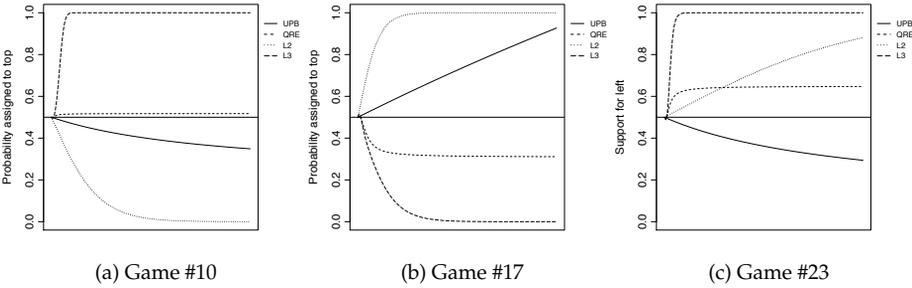


Figure 3.B.2: Concepts' predictions as function of the parameter value

We designed type *L2* games to distinguish between UPB, QRE, L3 and L2. Take game #2 as example and consider the concepts predictions in Figure 3.2.1b. We first discuss the L2 belief about the *Row Player*. Under L2 the *Column Player* believes that the *Row Player* believes that the *Column Player* plays both choices with equal probability. Thus the *Column Player* must believe that the *Row Player* plays *top*. For the L3 belief we have to go one step further in our analysis. Here the *Column Player* believes that the *Row Player* believes that the *Column Player* believes that the *Row Player* plays both choices with equal probability. Thus the *Column Player* must believe that the *Row Player* believes that the *Column Player* plays *left*. Therefore, the *Column Player* believes that the *Row Player* plays *bottom*. Under both UPB and QRE, the *Column Player* first believes that the *Row Player* believes that he plays both choices with equal probability. Hence, the *Column Player* deems it more likely that the *Row Player* will play *top*. Therefore his choice *left* becomes more attractive. If the *Column Player* now believes that the *Row Player* believes it is more likely that the *Column Player* will play *left*, he must believe that the *Row Player* is more likely to play *bottom*. Since the *Column Player* is not certain that the *Row Player* plays *bottom*, he still assigns substantial probability to him playing *top*, making his own choice *left* more attractive. At this point we enter a cycle where the *Column Player* believes the *Row Player* is more likely to play *bottom* and believes that the *Row Player* believes that he is more likely to play *left*.

L3 & QRE Games

L3 & QRE games are designed to differentiate between UPB, L2 and L3 and QRE. The idea behind these games is to test to what extent a player considers his own incentives when forming a belief about his opponent. Consider game #10 and the concepts' predictions in Figure 3.B.2a, if the *Column Player* believes that the *Row Player* believes that he will play both his choices with equal probability, then he must believe that the *Row Player* will play *bottom* rather than *top*, which is the L2 belief. However, if he reasons one step further, up to L3, he must believe that the *Row Player* believes that the *Column Player* believes that the *Row Player* plays both actions with equal probability. This means that he believes that the *Row Player* believes the *Column Player* will play *right*. Therefore, he must believe that the *Row Player* will play *top*. Since both UPB and QRE still assign enough probability to *Row Player* playing *bottom*, *right* will be slightly preferred for the *Column Player*. Under UPB the *Column Player* believes that the *Row Player* believes he will play both actions with almost equal probability and thus he believes that the *Row Player* will play *bottom*. However, under QRE the preference for *right* is slightly bigger and thus the *Column Player* thinks that the *Row Player* will play both choices with almost equal probability.

UPB vs QRE games

UPB vs QRE games are especially interesting because they separate UPB and QRE in 2×2 games without reducing QRE's predictions to .5 beliefs. As example we take game #17. The concepts predictions are depicted in Figure 3.B.2b illustrates this nicely. We consider the *Column Player's* belief about the *Row Player*. A UPB type starts by believing that his opponent believes that he will play both options with equal probability. Then *top* is more attractive than *bottom*. Thus the *Column Player* will put more weight on the *Row Player* playing *top*. Therefore he believes that the *Row Player* must deem his choice *right* more likely. However, only if the *Row Player* believes with substantial support that the *Column Player* plays *right*, he will choose *bottom*. So his choice *top* remains optimal and the cycle repeats. For QRE we again start by considering the *Column Player's* belief about the *Row Player*, when the *Column Player* believes that the *Row Player* believes that he plays both options with equal probability. Then the *Column Player* will deem the *Row Player's* choice *top* more likely than his choice *bottom*. Hence, the *Row Player* will deem *left* very unlikely. This is the difference between UPB

and QRE we discussed earlier, QRE reacts stronger to small differences in utilities than UPB. Hence, QRE deems *left* as very unlikely whereas UPB still assigns substantial probability to *left*. Since QRE assigns almost all weight to *right*, the *Row Player's* choice *bottom* becomes more likely than his choice *top*. Therefore, QRE predicts *bottom* compared to UPB which predicts *top*.²³ The belief formation of L2 and L3 follows exactly the same pattern as discussed in the two previous example.

UPB games

The games that uniquely identify UPB are 2×3 instead of 2×2 games. A dominated strategy is added to "trick away" L2 beliefs. Game #23 is such a game, where L2 deems the *Row Player's* choice *top* more likely even though essential support comes from the dominated choice *right* for the *Column Player*. Furthermore, L2 deems the *Column Player's* choice *left* more likely than *middle* and *right*. L3 best responds to these beliefs and therefore puts most the weight on *top* and *middle*. The reason why UPB and QRE's predictions are very similar to those of type 3 games. The concepts' predictions are depicted in Figure 3.B.2c.

3.C Appendix - Games

L2 Games

<p>Game #1</p> <table style="margin: auto; border-collapse: collapse;"> <tr> <td style="padding: 5px;"></td> <td style="padding: 5px; text-align: center;">L</td> <td style="padding: 5px; text-align: center;">R</td> </tr> <tr> <td style="padding: 5px; text-align: center;">T</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">337, 112</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">239, 141</td> </tr> <tr> <td style="padding: 5px; text-align: center;">B</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">165, 242</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">419, 27</td> </tr> </table>		L	R	T	337, 112	239, 141	B	165, 242	419, 27	<p>Game #2</p> <table style="margin: auto; border-collapse: collapse;"> <tr> <td style="padding: 5px;"></td> <td style="padding: 5px; text-align: center;">L</td> <td style="padding: 5px; text-align: center;">R</td> </tr> <tr> <td style="padding: 5px; text-align: center;">T</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">162, 423</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">441, 364</td> </tr> <tr> <td style="padding: 5px; text-align: center;">B</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">361, 305</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">236, 314</td> </tr> </table>		L	R	T	162, 423	441, 364	B	361, 305	236, 314	<p>Game #3</p> <table style="margin: auto; border-collapse: collapse;"> <tr> <td style="padding: 5px;"></td> <td style="padding: 5px; text-align: center;">L</td> <td style="padding: 5px; text-align: center;">R</td> </tr> <tr> <td style="padding: 5px; text-align: center;">T</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">191, 156</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">489, 70</td> </tr> <tr> <td style="padding: 5px; text-align: center;">B</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">432, 36</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">257, 488</td> </tr> </table>		L	R	T	191, 156	489, 70	B	432, 36	257, 488
	L	R																											
T	337, 112	239, 141																											
B	165, 242	419, 27																											
	L	R																											
T	162, 423	441, 364																											
B	361, 305	236, 314																											
	L	R																											
T	191, 156	489, 70																											
B	432, 36	257, 488																											
<p>Game #4</p> <table style="margin: auto; border-collapse: collapse;"> <tr> <td style="padding: 5px;"></td> <td style="padding: 5px; text-align: center;">L</td> <td style="padding: 5px; text-align: center;">R</td> </tr> <tr> <td style="padding: 5px; text-align: center;">T</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">342, 423</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">174, 496</td> </tr> <tr> <td style="padding: 5px; text-align: center;">B</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">34, 448</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">492, 273</td> </tr> </table>		L	R	T	342, 423	174, 496	B	34, 448	492, 273	<p>Game #5</p> <table style="margin: auto; border-collapse: collapse;"> <tr> <td style="padding: 5px;"></td> <td style="padding: 5px; text-align: center;">L</td> <td style="padding: 5px; text-align: center;">R</td> </tr> <tr> <td style="padding: 5px; text-align: center;">T</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">399, 319</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">360, 258</td> </tr> <tr> <td style="padding: 5px; text-align: center;">B</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">409, 94</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">55, 158</td> </tr> </table>		L	R	T	399, 319	360, 258	B	409, 94	55, 158	<p>Game #6</p> <table style="margin: auto; border-collapse: collapse;"> <tr> <td style="padding: 5px;"></td> <td style="padding: 5px; text-align: center;">L</td> <td style="padding: 5px; text-align: center;">R</td> </tr> <tr> <td style="padding: 5px; text-align: center;">T</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">376, 134</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">68, 246</td> </tr> <tr> <td style="padding: 5px; text-align: center;">B</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">344, 466</td> <td style="padding: 5px; border: 1px solid black; text-align: center;">289, 361</td> </tr> </table>		L	R	T	376, 134	68, 246	B	344, 466	289, 361
	L	R																											
T	342, 423	174, 496																											
B	34, 448	492, 273																											
	L	R																											
T	399, 319	360, 258																											
B	409, 94	55, 158																											
	L	R																											
T	376, 134	68, 246																											
B	344, 466	289, 361																											

²³Note that this type of reasoning is only possible for values of λ for which the logit map is still a contraction mapping.

CHAPTER 3. IDENTIFYING STRATEGIC THINKING IN ONE-SHOT GAMES

Game #7

	L	R
T	440, 378	339, 499
B	202, 482	348, 358

Game #8

	L	R
T	361, 232	76, 470
B	328, 294	367, 61

L3 & QRE Games

Game #9

	L	R
T	142, 497	146, 190
B	313, 20	36, 387

Game #10

	L	R
T	137, 244	341, 64
B	399, 204	96, 397

Game #11

	L	R
T	393, 13	208, 396
B	119, 327	435, 11

Game #12

	L	R
T	314, 477	452, 14
B	341, 66	401, 469

Game #13

	L	R
T	483, 342	225, 493
B	169, 387	494, 206

Game #14

	L	R
T	296, 5	184, 119
B	146, 373	354, 266

Game #15

	L	R
T	393, 345	64, 415
B	74, 457	365, 371

Game #16

	L	R
T	413, 159	198, 296
B	346, 401	278, 274

UPB v QRE Games

Game #17

	L	R
T	292, 94	35, 473
B	242, 351	47, 182

Game #18

	L	R
T	265, 37	270, 368
B	25, 367	355, 202

Game #19

	L	R
T	102, 452	329, 241
B	305, 158	234, 252

Game #20

	L	R
T	326, 491	459, 451
B	478, 36	4, 252

UPB Games

Game #21

	L	M	R
T	174, 150	74, 84	190, 447
B	94, 446	147, 13	208, 237

Game #22

	L	M	R
T	372, 447	135, 12	404, 17
B	338, 158	297, 7	239, 316

Game #23

	L	M	R
T	447, 238	295, 371	301, 192
B	381, 454	433, 208	223, 14

Game #24

	L	M	R
T	66, 461	388, 179	324, 451
B	229, 284	257, 62	281, 464

Game #25

	L	M	R
T	88, 349	161, 313	445, 30
B	182, 32	153, 459	354, 13

Game #26

	L	M	R
T	228, 4	301, 7	318, 302
B	348, 7	189, 309	327, 221

Game #27				Game #28			
	L	M	R		L	M	R
T	154, 75	160, 169	171, 424	T	163, 113	127, 466	196, 199
B	275, 115	33, 467	195, 440	B	277, 93	145, 122	72, 234

3.D Appendix - Robustness Classification

Table 3.6: Robustness Classification

Concept	No. of subjects		
	$\lambda = 2$	$\lambda = 3$	$\lambda = 4$
L1	4	4	4
L2	5	13	26
L3	16	20	28
UPB	2	2	6
Total	27	39	64

3.E Appendix - Instructions and Control Questions

Instructions for the Experiment

You are now taking part in an experiment involving economic decision-making.

During the experiment, you will be able to earn money in the form of "coins". At the end of the experiment, all the coins you earned will be converted into EUR and paid out at the following rate:

120 coins = 10 EUR

The amount you earn depends on your decisions and those of the other participants. At no time will the identity of the other participants be revealed. Neither will the other participants find out who you are.

All data and answers will be evaluated anonymously. You have drawn a personal cubicle number to guarantee your anonymity.

It is very important that you refrain from talking to the other participants at any time during the experiment. Please put away your pens, mobile phones, MP3 players, books, etc.

Read the instructions for the experiment carefully. If there is anything you don't understand during the experiment, take another look at the instructions. If you still have any questions, please raise your hand. We will then come to your cubicle and answer your questions in person.

The experiment consists of several parts. At the end, we would like you to answer a questionnaire.

First part of the experiment

You will be presented with 32 tables. The tables are like the one in this **example**:

		OTHER PLAYER	
		LEFT	RIGHT
YOU	TOP	Amount A1, Amount B1	Amount A2, Amount B2
	DOWN	Amount A3, Amount B3	Amount A4, Amount B4

In each table, you will play against another player. Both players will choose between 2 options. Half of the participants will choose between the options TOP and DOWN, the other half of the participants will choose between the options LEFT and RIGHT.

You can choose either TOP or DOWN.

You will be randomly and anonymously matched with another player (possibly the same, possibly a different player) who chooses between LEFT and RIGHT. For each table, the computer will randomly determine against which player you play in the respective table.

Note that you will NOT know ANY of the other player's decisions at ANY point during the experiment. Similarly, the other player will NOT know ANY of your decisions at ANY point during the experiment.

The payoffs of both players depend on the choices of both players. This means that your payoff depends both on your own choice and on the choice of the other player. Similarly, the payoff of the other player depends on his/her own choice and on your choice.

BOX 1

In particular,

- If you select TOP and the other player selects LEFT you will receive amount A1 and the other player will receive amount B1.
- If you select DOWN and the other player selects LEFT you will receive amount A3 and the other player will receive amount B3.
- If you select TOP and the other player selects RIGHT you will receive amount A2 and the other player will receive amount B2.
- If you select DOWN and the other player selects RIGHT you will receive amount A4 and the other player will receive amount B4.

The amounts in the tables can be between 0 and 500.

Task 1

For most of the tables, before making your choice you will be asked the following:

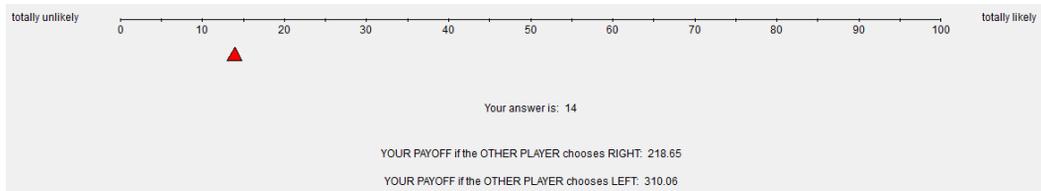
“How likely do you regard the event that the OTHER PLAYER chooses RIGHT?”

You will be asked to provide your answer on a scale from 0 to 100 by moving a red triangular cursor on the screen.

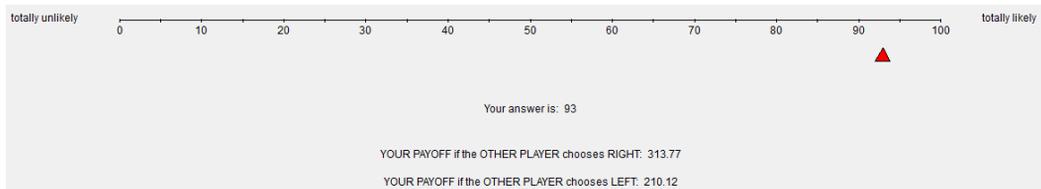
- “0” means that you regard the event that the other player chooses RIGHT totally unlikely (and thereby the event that the other player chooses LEFT totally likely)
- “100” means that you regard the event that the other player chooses RIGHT totally likely (and thereby the event that the other player chooses LEFT totally unlikely)

Each number from 0 to 100 assigns a payoff to you depending on the choice of the other player. The payoffs will appear once you first moved the cursor. Let's consider the following examples (see next page):

Example 1:



Example 2:



If you, like in Example 1, indicate 14, then your payoff is 218.65 coins if the other player chooses RIGHT and your payoff is 310.06 coins if the other player chooses LEFT. Similarly if you, like in Example 2, indicate 93, then your payoff is 313.77 coins if the other player chooses RIGHT and your payoff is 210.12 coins if the other player chooses LEFT.

For an overview, on the next page you are also given the complete list of payoffs that will result for each possible answer (BOX 2).

Note:

- The more certain you are that the other player will choose RIGHT, the closer to 100 your answer should be. This is because the closer your answer is to 100, the higher is your payoff in the case that the other player actually chooses RIGHT.
- The more certain you are that the other player will choose LEFT, the closer to 0 your answer should be. This is because the closer your answer is to 0, the higher is your payoff in the case that the other player actually chooses LEFT.

BOX 2

How likely do you regard the event that the OTHER PLAYER chooses RIGHT?

Your answer	Your payoff if the OTHER PLAYER chooses:		Your answer	Your payoff if the OTHER PLAYER chooses:	
	RIGHT	LEFT		RIGHT	LEFT
0	200,92	315,00	51	254,49	251,31
1	202,27	314,97	52	256,94	250,61
2	203,62	314,90	53	259,33	249,90
3	204,94	314,77	54	261,68	249,17
4	206,26	314,60	55	263,97	248,42
5	207,56	314,37	56	266,21	247,67
6	208,85	314,09	57	268,41	246,90
7	210,12	313,77	58	270,55	246,12
8	211,38	313,39	59	272,64	245,32
9	212,63	312,96	60	274,68	244,51
10	213,86	312,48	61	276,67	243,68
11	215,08	311,95	62	278,61	242,85
12	216,28	311,37	63	280,50	241,99
13	217,48	310,74	64	282,34	241,13
14	218,65	310,06	65	284,13	240,25
15	219,82	309,33	66	285,87	239,36
16	220,97	308,55	67	287,56	238,45
17	222,11	307,72	68	289,20	237,53
18	223,23	306,84	69	290,78	236,60
19	224,34	305,90	70	292,32	235,65
20	225,44	304,92	71	293,81	234,69
21	226,52	303,89	72	295,24	233,72
22	227,59	302,80	73	296,63	232,73
23	228,65	301,67	74	297,96	231,73
24	229,69	300,48	75	299,25	230,72
25	230,72	299,25	76	300,48	229,69
26	231,73	297,96	77	301,67	228,65
27	232,73	296,63	78	302,80	227,59
28	233,72	295,24	79	303,89	226,52
29	234,69	293,81	80	304,92	225,44
30	235,65	292,32	81	305,90	224,34
31	236,60	290,78	82	306,84	223,23
32	237,53	289,20	83	307,72	222,11
33	238,45	287,56	84	308,55	220,97
34	239,36	285,87	85	309,33	219,82
35	240,25	284,13	86	310,06	218,65
36	241,13	282,34	87	310,74	217,48
37	241,99	280,50	88	311,37	216,28
38	242,85	278,61	89	311,95	215,08
39	243,68	276,67	90	312,48	213,86
40	244,51	274,68	91	312,96	212,63
41	245,32	272,64	92	313,39	211,38
42	246,12	270,55	93	313,77	210,12
43	246,90	268,41	94	314,09	208,85
44	247,67	266,21	95	314,37	207,56
45	248,42	263,97	96	314,60	206,26
46	249,17	261,68	97	314,77	204,94
47	249,90	259,33	98	314,90	203,62
48	250,61	256,94	99	314,97	202,27
49	251,31	254,49	100	315,00	200,92
50	252,00	252,00			

Task 2

For each table, choose either TOP or DOWN.

How your payment is determined

It will be randomly determined whether you will be paid according to your decision in Task 1 or Task 2.

Payment if Task 1 has been selected

1 of your decisions in Task 1 will be randomly selected for payment. Your payment is determined by the number you indicated on the scale from 0 to 100 and the choice of the other player with whom you are matched in the table selected for payoff (see BOX 2).

Payment if Task 2 has been selected

1 of your decisions in Task 2 will be randomly selected for payment. Your payment is determined by your choice and the choice of the other player with whom you are matched in the table selected for payoff (see BOX 1).

You might ask yourself how we want you to act. For the scientific value of the experiment it is important that all participants set out to maximize their own individual payoffs. Therefore, we ask you to simply focus on getting as high payoffs as possible for yourself. All participants have been given this information asking them to maximize their own expected payoff.

Control Questions

Please answer the following questions:

Example 1

		OTHER PLAYER	
		LEFT	RIGHT
YOU	TOP	273, 91	24, 364
	DOWN	215, 301	32, 164

1. In Example 1,
 - A. If you select DOWN and the other player selects LEFT:
You will receive _____ coins and the other player will receive _____ coins.
 - B. If you select TOP and the other player selects RIGHT:
You will receive _____ coins and the other player will receive _____ coins.

Example 2

		OTHER PLAYER	
		LEFT	RIGHT
YOU	TOP	167, 185	487, 124
	MIDDLE	115, 78	26, 189
	DOWN	496, 250	368, 273

2. In Example 2,
 - A. If you select MIDDLE and the other player selects LEFT:
You will receive _____ coins and the other player will receive _____ coins.
 - B. If you select TOP and the other player selects LEFT:
You will receive _____ coins and the other player will receive _____ coins.
 - C. If you select MIDDLE and the other player selects RIGHT:
You will receive _____ coins and the other player will receive _____ coins.
 - D. If you select DOWN and the other player selects RIGHT:
You will receive _____ coins and the other player will receive _____ coins.

3. If your answer to the question "How likely do you regard the event that the OTHER PLAYER chooses RIGHT?" is 2:

A. Your payoff in case the other player chooses RIGHT is _____ coins.

B. Your payoff in case the other player chooses LEFT is _____ coins.

4. If your answer to the question "How likely do you regard the event that the OTHER PLAYER chooses RIGHT?" is 75:

A. Your payoff in case the other player chooses RIGHT is _____ coins.

B. Your payoff in case the other player chooses LEFT is _____ coins.

5. For each table,

you will always play against the same player.

you will always play against a different player and you will never play against the same player again.

you will play against a different player, but you might play against the same player more than once.

6. For 1 of the 32 tables you will get a payment in

Task 1 AND Task 2

Task 1 OR Task 2

Chapter 4

Local Prior Expected Utility: A Basis for Utility Representations Under Uncertainty

Ever since Ellsberg (1961) published his well-known paradox the matter of ambiguity received great interest in decision-making. Based on his paradox, many models have been conceived in an attempt to reconcile theory with the paradox. One way of avoiding the paradox is to relax the Independence Axiom of Anscombe and Aumann (1963) (AA). This approach yielded two of the most well-known models of ambiguity: Schmeidler's (1989) Choquet Expected Utility (CEU), and Gilboa and Schmeidler's (1989) Maximin Expected Utility (MEU). Under both approaches the decision maker acts as if he maximizes utility with respect to a set of priors. In this paper we identify necessary assumptions on the preference relation to obtain a representation under which the decision maker maximizes utility with respect to a set of priors, called Local Prior Expected Utility (LEU). Moreover, we show that the prior used to evaluate a certain act is equal to the gradient of some appropriately defined utility mapping. We argue that the equality is not a mere technicality but coherent with the qualitative interpretation of a probabilistic belief. Based on this result we provide a unified

approach to MEU, CEU, and AA and characterize the respective sets of priors.¹

A preference relation \succeq on acts that satisfies the standard axioms Weak Order, Monotonicity, Continuity, Risk Independence, and Non-Degeneracy can be represented by a functional ψ that maps vectors of statewise utilities to "overall" utilities. We show that Monotonicity of ψ guarantees the existence of its Gâteaux derivative almost everywhere relying on a deep mathematical theorem by Chabrilac and Crouzeix (1987). Using only the Gâteaux derivative of ψ as an analytical tool, we prove characterization results corresponding to LEU, CEU, MEU, and AA. Thus, our proofs do not require an understanding of complicated mathematics and can be understood with basic knowledge of real analysis. Within our approach we clearly identify the structure of the set of priors needed in these environments.

As a basis for our approach, we identify an axiom called Independence of Certainty Equivalents that is weaker than Gilboa and Schmeidler's (1989) C-Independence and weaker than Schmeidler's (1989) Comonotonic Independence, but together with Weak Order, Monotonicity, Continuity, Risk Independence, Non-Degeneracy, and Uncertainty Aversion induces the same restrictions on the preference relation on acts as the Gilboa-Schmeidler axioms. We then show that without assuming Uncertainty Aversion we can represent a decision maker's preferences by taking expectations of an affine utility function u with respect to a (possibly) different prior for every act, which we call LEU. The prior used is equal to the gradient of ψ at the vector of utilities induced by the act if the Gâteaux derivative exists. In case the Gâteaux derivative does not exist at a particular vector, we can approximate it by the Gâteaux derivatives of nearby acts. This is possible since we show that the Gâteaux derivative exists almost everywhere, and the utility representation of the preference relation must be continuous. Thus non-smoothness of ψ does not affect our analysis.

Relaxing the Independence Axiom further by requiring invariance with respect to translations but not invariance to rescaling still results in a prior representation. However, the prior is no longer necessarily equal to the gradient of ψ at the vector of utilities induced by the act, but to the gradient at a possibly different vector. This result can be shown under Weak C-Independence in the sense of Maccheroni et al. (2006). If on the other hand, the invariance of translations of utility profiles is relaxed, there is no representation as in the previous two cases

¹This chapter is based on the paper "Local Prior Expected Utility: a basis for utility representations under uncertainty" which I wrote together with Andrés Perea.

since the gradient is not guaranteed to be a prior anymore.

The paper proceeds by introducing the fundamental axioms and necessary notation in Section 4.1. These axioms on preference relations over acts guarantee the existence of a function ψ on vectors of utilities, which represents the preference relation. Section 4.2 gives a formal definition of the Gâteaux derivative of ψ , shows that it exists almost everywhere and is equivalent to ψ 's gradient if taken into the direction of the location. Using the existence of the Gâteaux derivative almost everywhere, we show in Section 4.3 that under Weak Order, Monotonicity, Continuity, Risk Independence, Non-Degeneracy, and Independence of Certainty Equivalents ψ is C-Additive and homogeneous, which is sufficient to prove the representation result for LEU. In Section 4.4 we use this framework to prove the representation results corresponding to MEU, CEU, and the AA representation by showing how additional assumptions on the preference relation translate into features of ψ and its Gâteaux derivative. Thereby, we characterize each set of axioms by regions of acts that use the same prior. Finally, Section 3.4 discusses the implications of further relaxations of the Independence Axiom and the connection to Ghirardato et al. (2004) who also characterize the set of priors under C-Independence.

4.1 A Utility Function for Acts

We take the classical Anscombe-Aumann setup (Anscombe and Aumann, 1963) as a starting point, denote by $\Omega := \{1, \dots, n\}$ the finite set of states of nature and by X the finite set of outcomes. Then let $\Delta(X)$ be the set of all probability distributions over X . An act is a mapping $f : \Omega \rightarrow \Delta(X)$, and \mathcal{F} is the set of all such acts. By $f(\omega)$ we denote the probability distribution over X induced by f in state ω , and by $f(\omega)(x)$ the probability of receiving alternative x in state ω . If we take any lottery $p \in \Delta(X)$ and write $\langle p \rangle$, we mean the constant act such that $\langle p \rangle(\omega) = p$ for all $\omega \in \Omega$. Moreover, we define mixtures of acts point wise, i.e. for each $\alpha \in [0, 1]$, $\alpha f + (1 - \alpha)g$ is the act that gives the prize $\alpha f(\omega) + (1 - \alpha)g(\omega)$ in state $\omega \in \Omega$. Finally, for some affine function $u : \Delta(X) \rightarrow \mathbb{R}$ we define $u \circ f := (u(f(\omega)))_{\omega \in \Omega}$ and let $\mathbb{1}_\omega$ be the vector for $\omega \in \Omega$ where $\mathbb{1}_\omega(\omega) = 1$ and $\mathbb{1}_\omega(\omega') = 0$ for all $\omega' \in \Omega$ such that $\omega' \neq \omega$. By $\mathbb{1}_\Omega$ we then denote the vector with $\mathbb{1}_\Omega(\omega) = 1$ for all $\omega \in \Omega$. The following five Axioms are the basis for most models in decision theory.

Axiom 1 (Weak Order). For all $f, g, h \in \mathcal{F}$ we have either $f \succeq g$ or $g \succeq f$, and if $f \succeq g$ and $g \succeq h$, then $f \succeq h$.

Axiom 2 (Monotonicity). For any two acts $f, g \in \mathcal{F}$ with $\langle f(\omega) \rangle \succeq \langle g(\omega) \rangle$ for all $\omega \in \Omega$ we have $f \succeq g$.

Axiom 3 (Continuity). For any acts $f, g, h \in \mathcal{F}$ such that $f \succ g \succ h$, there exists some $\alpha \in (0, 1)$ such that $g \sim \alpha f + (1 - \alpha)h$.

Axiom 4 (Nondegeneracy). There exist $f, g \in \mathcal{F}$ such that $f \succ g$.

Axiom 5 (Risk Independence). If $x, y, z \in \Delta(X)$ and $\lambda \in (0, 1]$ then $\langle x \rangle \succ \langle y \rangle$ implies $\lambda \langle x \rangle + (1 - \lambda) \langle z \rangle \succ \lambda \langle y \rangle + (1 - \lambda) \langle z \rangle$.

Proposition 4.1 is a well-known result in the literature (e.g. Cerreia-Vioglio et al., 2011, Ok, 2007). For the sake of completeness we reproduced the proof based on Ok (2007) in Appendix 2.A.

Proposition 4.1. *If the preference relation \succeq satisfies Axioms 1 to 5, then there exists an affine function u on $\Delta(X)$ with $u(p) \in [-1, 1]$ for all $p \in \Delta(X)$, $u(p^*) = 1$ for some $p^* \in \Delta(X)$ and $u(p_*) = -1$ for some $p_* \in \Delta(X)$, a utility function U on \mathcal{F} with $U(\langle p \rangle) = u(p)$ for all $p \in \Delta(X)$, and a continuous and monotone function $\psi : [-1, 1]^\Omega \rightarrow \mathbb{R}$ such that*

$$U(f) = \psi(u \circ f),$$

for all $f \in \mathcal{F}$. Moreover, for every act $f \in \mathcal{F}$ we can find a constant act $p \in \Delta(X)$ such that $f \sim \langle p \rangle$.

Note that it follows from the definition of ψ and u that $\psi(\mathbb{1}_\Omega) = \psi(u \circ \langle p^* \rangle) = U(\langle p^* \rangle) = 1$. Thus, we can evaluate each act $f \in \mathcal{F}$ by computing a vector of utilities $u \circ f$ using the utility function u on lotteries, which we can then evaluate using ψ , a monotone and continuous function. The function ψ simplifies the mathematical treatment since we can directly work with utility vectors rather than vectors of lotteries.

4.2 Gradient of ψ

In the following we will establish that the Gâteaux derivative, a generalized directional derivative, of ψ exists almost everywhere. This is due to the fact that

ψ is monotonic, i.e. for any $F, G \in [-1, 1]^\Omega$ we have $\psi(F) \leq \psi(G)$ whenever $F(\omega) \leq G(\omega)$ for all $\omega \in \Omega$. Formally, the Gâteaux derivative is defined as follows.

Definition 4.1 (Gâteaux derivative). The function ψ is said to be Gâteaux differentiable at $F \in [-1, 1]^\Omega$ if

1. $\psi'(F, G) := \lim_{\alpha \downarrow 0} \frac{\psi(F + \alpha G) - \psi(F)}{\alpha}$, exists for all $G \in [-1, 1]^\Omega$, and if
2. $\psi'(F, \lambda G + \gamma H) = \lambda \psi'(F, G) + \gamma \psi'(F, H)$ for all $G, H \in [-1, 1]^\Omega$ and all $\lambda, \gamma \in \mathbb{R}$ such that $\lambda G + \gamma H \in [-1, 1]^\Omega$.

However, it is not guaranteed that the Gâteaux derivative exists everywhere. By combining Theorem 14 in Chabrilac and Crouzeix (1987) and the fact that Fréchet differentiability implies Gâteaux differentiability, we obtain the following Theorem.

Theorem 4.1. *Monotone functions on \mathbb{R}^n are almost everywhere Gâteaux differentiable.*

It ensures the existence of derivatives of monotone functions almost everywhere, i.e. in a set of points with measure one. Since ψ is monotonic we are thereby ensured that its Gâteaux derivative exist almost everywhere and thus the following corollary is immediate.

Corollary 4.1. *The function $\psi : [-1, 1]^\Omega \rightarrow \mathbb{R}$ is almost everywhere Gâteaux differentiable.*

The fact that ψ is not differentiable everywhere does not cause a problem. Later we will show how one can remedy this limitation. We proceed by defining the gradient of ψ . Recall that $\Omega = \{1, \dots, n\}$.

Definition 4.2. Let the gradient of ψ at $F \in [-1, 1]^\Omega$ be

$$(\nabla \psi)(F) := \left(\psi'(F, \mathbb{1}_1), \psi'(F, \mathbb{1}_2), \dots, \psi'(F, \mathbb{1}_n) \right),$$

where $\mathbb{1}_m \in [-1, 1]^\Omega$ for all $m \in \Omega$.

Instead of taking the derivative in the direction of every dimension, the Gâteaux derivative describes changes in multiple directions at the same time. Lemma 4.1 establishes the exact relation between the Gâteaux derivative and the gradient in our setting.

Lemma 4.1. *If the Gâteaux derivative of ψ at $F \in [-1, 1]^\Omega$ exists, it holds that $\nabla\psi(F)G = \psi'(F, G)$ for any $G \in [-1, 1]^\Omega$.*

Hence, taking the gradient of ψ at some F and computing the vector product with some G is equivalent to taking the Gâteaux derivative at F in the direction of G , where $F, G \in [-1, 1]^\Omega$. This gives us a simple way to go back and forth between the gradient and the Gâteaux derivative at some point. As we will see shortly this connection provides a simple way of not only proving a new representation result but also to prove the classical results in the ambiguity literature.

4.3 Representation for Local Prior Expected Utility

In order to obtain our representation result, we need to assume that the convex combination of an act $f \in \mathcal{F}$ with a constant act $\langle q \rangle$ is equivalent to the convex combination of f 's certainty equivalent $\langle p \rangle$ with the same constant act $\langle q \rangle$, which we formalize in the following axiom.

Axiom 6 (Independence of Certainty Equivalent). *If $f \in \mathcal{F}$ and $p, q \in \Delta(X)$ such that $f \sim \langle p \rangle$, then $\lambda f + (1 - \lambda)\langle q \rangle \sim \lambda\langle p \rangle + (1 - \lambda)\langle q \rangle$ for all $\lambda \in (0, 1]$.*

From now on we assume that \succeq also satisfies Axiom 6. This additional assumption gives rise to the following representation theorem.

Theorem 4.2. *Let \succeq be a preference relation on \mathcal{F} satisfying Axioms 1 to 6, represented by the functions ψ and u as described in Proposition 4.1, and suppose the Gâteaux derivative of ψ at $F \in [-1, 1]^\Omega$ exists, where $F = u \circ f$ for some $f \in \mathcal{F}$, then*

$$U(f) = \nabla\psi(F)F.$$

Moreover, $\nabla\psi(F)$ is a probability measure.

Hence, a decision maker satisfying Axioms 1 to 6 behaves as if he evaluates an act $f \in \mathcal{F}$ by computing the expected utility of f given the respective prior $\nabla\psi(F)$ and the utility function u . Therefore, our representation result is similar to the classical Anscombe-Aumann representation. Only instead of a single prior, the decision maker uses a possibly different prior for every act. Moreover, as in the Anscombe-Aumann setup, the prior is equal to the gradient at the respective act, with the difference that in their setup the gradient is constant over all acts.

That the gradient is considered to be a prior is not only due to the fact that it satisfies the characteristics of a probability measure. The interpretation of changes in the overall utility due to a small local deviation in the direction of a particular state is tightly connected to the meaning of a probability distribution. Since the gradient describes reactions in overall utility to small changes in the utility of a certain state, the value of the gradient must describe how likely the decision maker deems that particular state. Thus, the value of the gradient also qualitatively identifies the likelihood of states according to the decision maker's preference.

To prove the representation theorem we will make heavy use of two properties of ψ called *positive homogeneity* and *C-Additivity*, which we establish here.

Proposition 4.2. *If the preference relation \succeq satisfies Axioms 1 to 6, then ψ is*

1. *positively homogeneous, i.e. it holds that $\psi(\lambda F) = \lambda\psi(F)$ for all $\lambda > 0$ and all $F \in [-1, 1]^\Omega$ such that $\lambda F \in [-1, 1]^\Omega$, and it is*
2. *C-Additive, i.e. it holds that $\psi(F + \alpha \mathbb{1}_\Omega) = \psi(F) + \alpha\psi(\mathbb{1}_\Omega)$ for all $F \in [-1, 1]^\Omega$ and $\alpha \geq 0$ such that $F + \alpha \mathbb{1}_\Omega \in [-1, 1]^\Omega$.*

The existence of the Gâteaux derivative almost everywhere is such a strong result that we do not need to rely on any other advanced theorems. Positive homogeneity and C-Additivity of ψ together with the definition of the Gâteaux derivative are sufficient to prove the representation theorem. Lemma 4.2 is an important step as it establishes the connection between the utility of an act $f \in \mathcal{F}$ and the Gâteaux derivative at F in the direction of F if the Gâteaux derivative at F exists and $F = u \circ f$.

Lemma 4.2. *If the preference relation \succeq on \mathcal{F} satisfies Axioms 1 to 6, and the Gâteaux derivative of ψ at F exists, it holds that $\psi(F) = \psi'(F, F)$ for any $F \in [-1, 1]^\Omega$.*

When we combine Lemma 4.2 with Lemma 4.1, we can conclude that if the Gâteaux derivative of ψ exists at $F \in [-1, 1]^\Omega$, then $\psi(F) = \nabla\psi(F)F$. Thus, we are left with the task to show that $\nabla\psi(F)$ is a prior, which we establish in Proposition 4.3.

Proposition 4.3. *If the preference relation \succeq on \mathcal{F} satisfies Axioms 1 to 6 and the Gâteaux derivative of ψ at $F \in [-1, 1]^\Omega$ exists, the gradient $\nabla\psi(F)$ of ψ has the following properties:*

- $\nabla\psi(F)(\omega) \geq 0$ for all $\omega \in \Omega$, and
- $\sum_{\omega \in \Omega} \nabla\psi(F)(\omega) = 1$.

These are all the ingredients needed to prove Theorem 4.2. We show the formal proof here as it nicely illustrates how the previous results establish Theorem 4.2.

Proof of Theorem 4.2. It follows directly from Proposition 4.3 that for any $F \in [-1, 1]^\Omega$ the gradient $\nabla\psi(F)$ is a probability measure. By Lemma 4.1 and 4.2 we can now compute $\psi(F)$ using $\nabla\psi(F)F$. Since we defined ψ such that $U(f) = \psi(F)$, it follows that $U(f) = \nabla\psi(F)F$. \square

In Theorem 4.2 we have shown that Axioms 1 to 6 imply the Local Prior Expected Utility representation as given in the statement of that theorem. We will now provide, by means of Theorem 4.3, a counterpart to this result by showing that every Local Prior Expected Utility representation with certain properties implies the Axioms 1 to 6.

Theorem 4.3. *The preference relation \succeq satisfies Axioms 1 to 6 if and only if there exists*

- an affine function $u : \Delta(X) \rightarrow \mathbb{R}$ with $u(p) \in [-1, 1]$ for all $p \in \Delta(X)$, and with $u(p^*) = 1$ for some $p^* \in \Delta(X)$ and $u(p_*) = -1$ for some $p_* \in \Delta(X)$,
- a utility function $U : \mathcal{F} \rightarrow \mathbb{R}$ representing \succeq such that $U(\langle p \rangle) = u(p)$ for all $p \in \Delta(X)$, and such that for every $f \in \mathcal{F}$ there is a $q \in \Delta(X)$ with $U(f) = U(\langle q \rangle)$, and
- a continuous, monotone, positively homogeneous, and C-Additive function $\psi : [-1, 1]^\Omega \rightarrow \mathbb{R}$ such that $U(f) = \psi(u \circ f) = \nabla\psi(F)F$ for every $f \in \mathcal{F}$ and $F = u \circ f$.

Finally, we discuss the situation when the Gâteaux derivative at $F \in [-1, 1]^\Omega$ does not exist. By continuity of ψ we can simply approximate the utility at such a point. Thus, from here on we assume that whenever the derivative for any particular act does not exist, we approximate it as suggested in Proposition 4.4.

Proposition 4.4. *If the preference relation \succeq on \mathcal{F} satisfies Axioms 1 to 6 and the Gâteaux derivative at $F \in [-1, 1]^\Omega$ does not exist, then there is a sequence $F_n \rightarrow F$ such that the Gâteaux derivative exists at every F_n , and such that $\psi(F) = \lim_{n \rightarrow \infty} \nabla\psi(F_n)F_n$.*

4.4 Representation for Classical Models of Decision Making Under Uncertainty

4.4.1 Maximin Expected Utility

Gilboa and Schmeidler (1989) relax the Full-Independence axiom of Anscombe and Aumann (1963) and replace it with a much weaker Axiom called C-Independence.

Axiom 7. For any $f, g \in \mathcal{F}$, $p \in \Delta(X)$, and any $\lambda \in (0, 1]$, $f \succeq g$ if and only if $\lambda f + (1 - \lambda)\langle p \rangle \succeq \lambda g + (1 - \lambda)\langle p \rangle$.

It is a weakening of the Full-Independence Axiom in the sense that it only requires the independence property with respect to mixtures with constant acts. However, it turns out that we can prove Gilboa and Schmeidler's (1989) representation result with a weaker set of axioms. We do not need to assume Axiom 7 but can simply continue to work with Axioms 5 and 6. To see why those are weaker note that C-Independence implies both Axiom 5 and Axiom 6. Firstly, if we take two constant acts $\langle f \rangle, \langle g \rangle \in \mathcal{F}$ then Axiom 7 directly implies Axiom 5. Secondly, if we take any $f \in \mathcal{F}$ and its certainty equivalent $\langle p \rangle$ then Axiom 7 implies Axiom 6. Furthermore note that a preference relation that satisfies Axioms 1 to 6 is C-Additive and positively homogeneous. Using these properties it can be shown that such a preference relation also satisfies Axiom 7. The reason we chose to work with Axioms 5 and 6 instead of Axiom 7 is that this approach pinpoints the exact properties we need to prove the representation theorem.

The C-Independence Axiom was not the only novelty in Gilboa and Schmeidler's approach. They also introduced an axiom called Uncertainty Aversion, which implies that if the decision maker is indifferent between two acts then he will never prefer one of them to a mixture of the two acts.

Axiom 8 (Uncertainty Aversion). For any $f, g \in \mathcal{F}$, if $f \sim g$ then $\frac{1}{2}f + \frac{1}{2}g \succeq f$.

Under the assumption that \succeq satisfies Axioms 1 to 6 and Axiom 8 we can now prove their famous representation theorem. Note that our proof is again very simple in nature as it only exploits the definition of the Gâteaux derivative. In order to state their representation theorem we define

$$\mathcal{P}_G := \left\{ \nabla \psi(F) : \text{Gâteaux derivative of } \psi \text{ at } F \in [-1, 1]^\Omega \text{ exists} \right\}$$

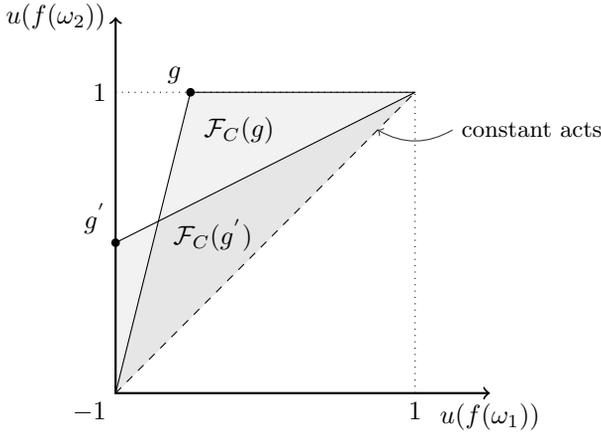


Figure 4.4.1: Acts represented as utility vectors for two states.

and $\mathcal{P} := \overline{\text{co}}(\mathcal{P}_G)$, where $\overline{\text{co}}$ denotes the convex closure.

Theorem 4.4. *A preference relation \succeq that satisfies Axioms 1 to 6 and Axiom 8 can be represented by*

$$U(f) := \min_{p \in \mathcal{P}} pF,$$

for any $f \in \mathcal{F}$ and $F = u \circ f$.

The theorem states that a decision maker with a preference relation \succeq , which satisfies Axioms 1 to 6 and Axiom 8, acts as if he evaluates each act $f \in \mathcal{F}$ using an affine utility function u on $\Delta(X)$ and the prior $p \in \mathcal{P}$ that induces the minimal expected utility for act f . The fact that the decision maker uses the prior that induces the minimal expected utility clearly stems from Axiom 8. Without Axiom 8 the decision maker evaluates each act with its respective prior, the gradient at that act, as shown in Theorem 4.2. Technically Axiom 8 causes ψ to be Super Additive, which we establish in the following proposition.

Proposition 4.5 (Super Additive). *Suppose that the preference relation \succeq on \mathcal{F} satisfies Axioms 1 to 6 and Axiom 8. Let \succeq be represented by ψ and u , as in Proposition 4.1, then for any $F, G \in [-1, 1]^\Omega$ with $F + G \in [-1, 1]^\Omega$ we have that $\psi(F + G) \geq \psi(F) + \psi(G)$.²*

²Proposition 4.5 is well-known. We reproduce the proof of Ok (2007) for the sake of completeness.

The proof of Theorem 4.4 shows directly how Super Additivity of ψ lets the decision maker evaluate acts using the prior that induces the minimal expected utility.

Proof of Theorem 4.4. Take $f \in \mathcal{F}$ and let $F := u \circ f$, then we have to show that $U(f) \leq \nabla\psi(G)F$ for all $\nabla\psi(G) \in \mathcal{P}_G$. This is sufficient since Theorem 4.2 still holds and we therefore know that $U(f) = \nabla\psi(F)F$ and that $\nabla\psi(G)$ is a probability measure for all $G \in [-1, 1]^\Omega$ whenever the Gâteaux derivative exists at G . By Lemma 4.1 it follows that $\nabla\psi(G)F = \psi'(G, F)$. Using the definition of the Gâteaux derivative

$$\begin{aligned} \psi'(G, F) &= \lim_{\alpha \downarrow 0} \frac{\psi(G + \alpha F) - \psi(G)}{\alpha} \\ &\geq \lim_{\alpha \downarrow 0} \frac{\psi(G) + \psi(\alpha F) - \psi(G)}{\alpha} \\ &= \lim_{\alpha \downarrow 0} \frac{\alpha\psi(F)}{\alpha} \\ &= \psi(F), \end{aligned}$$

where the inequality follows by Proposition 4.5, and the second equation by positive homogeneity of ψ . By definition $\psi(F) = U(f)$, and thus $U(f) \leq \nabla\psi(G)F$. Therefore we can evaluate every act using the prior $p \in \mathcal{P}$ that induces the minimum utility for act f . \square

It is natural to ask how the set of priors in Theorem 4.4 is structured. To answer this question, we define the set of all acts in the same C -independent region as $g \in \mathcal{F}$ by $\mathcal{F}_C(g) := \{\lambda g + (1 - \lambda)\langle p \rangle : \lambda \in (0, 1] \text{ and } p \in \Delta(X)\}$. Note that our definition of $\mathcal{F}_C(g)$ does not include constant acts. This fact, however, does not limit the scope of our description as the utility of constant acts is independent of the prior. The following Proposition shows that each set $\mathcal{F}_C(g)$ can be described by a single prior $p \in \mathcal{P}$.

Proposition 4.6. *For an act $g \in \mathcal{F}$ and all $f \in \mathcal{F}_C(g)$ it holds that $U(f) = \nabla\psi(G)F$, where $F = u \circ f$ and $G = u \circ g$.*

Proof. Fix some $g \in \mathcal{F}$ and take any $f \in \mathcal{F}_C(g)$. Thus we must have for some $\lambda \in (0, 1]$ and some $p \in \Delta(X)$ that $f = \lambda g + (1 - \lambda)\langle p \rangle$. Let $u(p) = \beta$. Then by

affinity of u we have

$$\begin{aligned}
 F &= u \circ (\lambda g + (1 - \lambda)\langle p \rangle) \\
 &= \lambda u \circ g + (1 - \lambda)u \circ \langle p \rangle \\
 &= \lambda G + (1 - \lambda)\beta \mathbb{1}_\Omega.
 \end{aligned}$$

Then,

$$\begin{aligned}
 \nabla\psi(G)F &= \psi'(G, F) \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi(G + \alpha F) - \psi(G)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi(G + \alpha(\lambda G + (1 - \lambda)\beta \mathbb{1}_\Omega)) - \psi(G)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi((1 + \alpha\lambda)G + \alpha(1 - \lambda)\beta \mathbb{1}_\Omega) - \psi(G)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{(1 + \alpha\lambda)\psi(G) + \alpha(1 - \lambda)\beta - \psi(G)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{\alpha\lambda\psi(G) + \alpha(1 - \lambda)\beta}{\alpha} \\
 &= \lambda\psi(G) + (1 - \lambda)\beta \\
 &= \psi(\lambda G + (1 - \lambda)\beta \mathbb{1}_\Omega) \\
 &= \psi(F) \\
 &= U(f),
 \end{aligned}$$

where the fifth and the eighth equality hold by C-Additivity and positive homogeneity of ψ . □

To understand how the set of priors \mathcal{P} looks like, observe that there could be $g, g' \in \mathcal{F}$ with $g \neq g'$ such that $\mathcal{F}_C(g) \cap \mathcal{F}_C(g') \neq \emptyset$ as depicted in Figure 4.4.1.³ Thus we define the set $\mathcal{G} := \{g \in \mathcal{F} : g \notin \mathcal{F}_C(g') \text{ for any } g' \neq g\}$ as the set of acts that induce the "largest" C-independent region. For the two-state case we have $\mathcal{G} = \{(-1, 1), (1, -1)\}$ and therefore we only need two priors to describe the decision makers' preferences. Figure 4.4.2 gives an example of a $\mathcal{F}_C(g)$ with $g \in \mathcal{G}$ for the three-state case. Thus the set of priors can be exactly described by

³For a matter of presentation we simplified the axes' labels. Complete labels would read: $u(f(\omega_i)), u(g(\omega_i))$, and $u(g'(\omega_i))$ for state $i \in \Omega$.

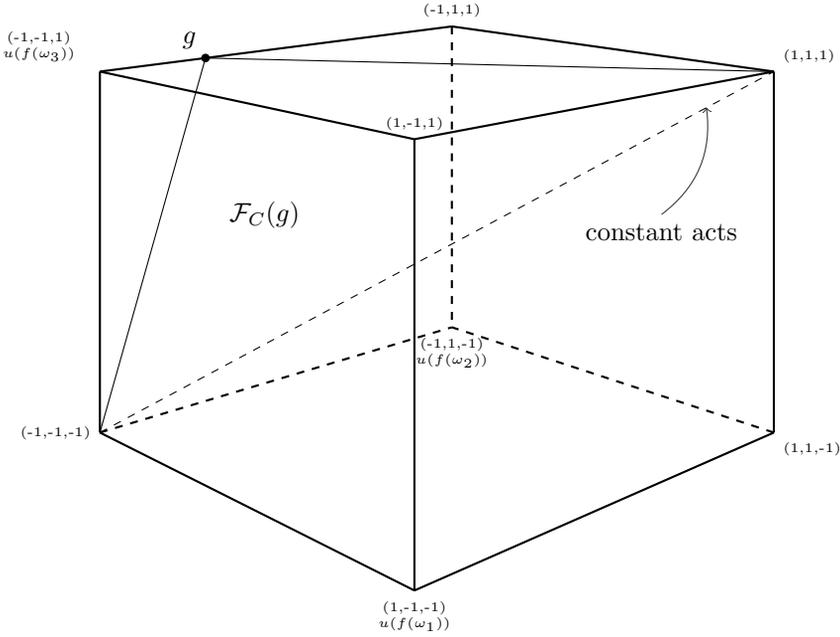


Figure 4.4.2: Acts represented as utility vectors for three states.

$\mathcal{P}_G = \{\nabla\psi(G) : G = u \circ g \text{ with } g \in \mathcal{G}\}$. Since the set \mathcal{G} is possibly infinite when there are more than 2 states, the set \mathcal{P}_G could be as well.

4.4.2 Choquet Expected Utility

Schmeidler's (1989) Choquet Expected Utility also relaxes the Independence Axiom of Anscombe-Aumann. Schmeidler introduces the notion of Comonotonicity and requires independence only for acts that are pairwise comonotonic. The idea behind Comonotonicity is that two comonotonic acts induce the same ranking over states.

Definition 4.3 (Comonotonicity). Two acts $f, g \in \mathcal{F}$ are comonotonic if for no $\omega, \omega' \in \Omega$, it holds that $\langle f(\omega) \rangle \succ \langle f(\omega') \rangle$ and $\langle g(\omega') \rangle \succ \langle g(\omega) \rangle$.

Building on the definition of Comonotonicity we can introduce Schmeidler's axiom of Comonotonic Independence.

Axiom 9 (Comonotonic Independence). For all pairwise comonotonic acts $f, g, h \in \mathcal{F}$ and for all $\lambda \in (0, 1)$, if $f \succ g$ then $\lambda f + (1 - \lambda)h \succ \lambda g + (1 - \lambda)h$.

Comonotonic Independence is stronger than Independence of Certainty Equivalent and Risk Independence. However, we cannot compare it with C-Independence because Comonotonic Independence requires all three acts to be pairwise comonotonic. C-Independence, on the other hand, only requires the mixing act to be constant and thus Comonotonic Independence does not imply C-Independence. To see why Comonotonic Independence implies Risk Independence, take all three acts to be constant acts.⁴ Since these three constant acts are pairwise comonotonic, Comonotonic Independence implies Risk Independence. For Independence of Certainty Equivalents we take an act, its certainty equivalent, and a constant act. Again, all acts are pairwise comonotonic because the certainty equivalent is also a constant act. Thus we can conclude that Comonotonic Independence implies Independence of Certainty Equivalents, which also illustrates why Comonotonic Independence leads to a stronger additivity result.

Proposition 4.7 (Comonotonic Additivity). *Assume that \succeq satisfies Axioms 1 to 4 and Axiom 9. Let \succeq be represented by ψ and u as defined in Proposition 4.1, and take any two comonotonic acts $f, g \in \mathcal{F}$ and $F, G \in [-1, 1]^\Omega$ such that $u \circ f = F$ and $u \circ g = G$. Then it holds that $\psi(F + \alpha G) = \psi(F) + \alpha\psi(G)$ for any $\alpha > 0$ such that $F + \alpha G \in [-1, 1]^\Omega$.*

Compared to C-Additivity, Comonotonic Additivity also holds for two comonotonic acts and not only for an act and a constant act. It is immediate that Comonotonic Additivity implies C-Additivity. To discuss Choquet Expected Utility we need to introduce some further notation and the notion of a capacity. A capacity is a set valued function that assigns a number between 0 and 1 to every subset of Ω . Furthermore, it must hold that a subset cannot receive a higher value than its superset.

Definition 4.4 (Capacity). A function $v : 2^\Omega \rightarrow \mathbb{R}$ is a capacity if it satisfies the following two conditions:

- $v(\emptyset) = 0$ and $v(\Omega) = 1$, and
- for all $E, G \in 2^\Omega$, we have that $E \subset G$ implies $v(E) \leq v(G)$.

Moreover, recall that $n := |\Omega|$, fix an $F \in [-1, 1]^\Omega$ and let $\sigma : \{1, \dots, n\} \rightarrow \Omega$ be a bijective function such that we have $F(\sigma(1)) \geq F(\sigma(2)) \geq \dots \geq F(\sigma(n))$.

⁴Proposition 4.20 in Appendix 2.D establishes that given Axioms 1 to 4 and Axiom 9 it holds that for all pairwise comonotonic acts $f, g, h \in \mathcal{F}$ such that $f \sim g$ we have $\lambda f + (1 - \lambda)h \sim \lambda g + (1 - \lambda)h$ where $\lambda \in (0, 1)$.

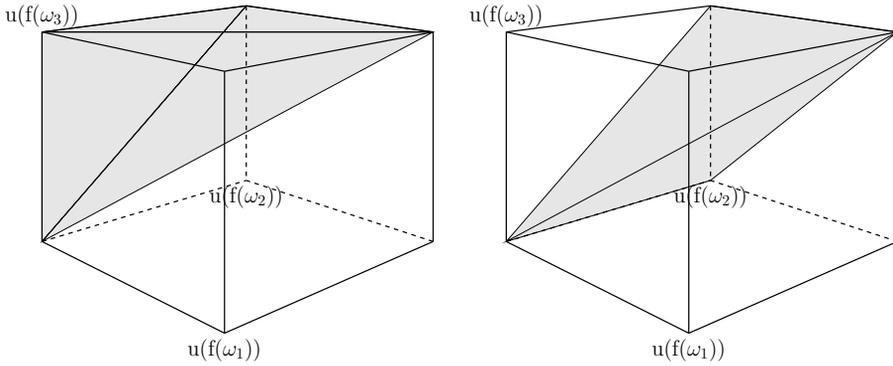


Figure 4.4.3: Example of comonotonic regions for three states.

Moreover, we define $F(\sigma(n+1)) = 0$. Then the Choquet integral of F with respect to the capacity v is defined as

$$\int_{\Omega} F dv = \sum_{i=1}^n [F(\sigma(i)) - F(\sigma(i+1))] v(\{\sigma(1), \dots, \sigma(i)\}). \quad (4.4.1)$$

The idea behind Choquet Integration is to compute the area under a functions' graph by placing rectangles within the graph of the function. However, this is not done vertically as with Riemann integration but horizontally. An often used analogy is that of pouring water, i.e. filling up the graph from the bottom to the top with the value of the capacity as basis.⁵

These are all the ingredients needed to state and prove Schmeidler's representation theorem for Choquet Expected Utility.

Theorem 4.5. *For a preference relation \succeq that satisfies Axioms 1 to 4 and Axiom 9 there exists a unique capacity v on Ω and an affine real valued function u on $\Delta(X)$ such that*

$$U(f) = \int_{\Omega} (u \circ f) dv, \quad (4.4.2)$$

for any $f \in \mathcal{F}$.

So if a decision maker satisfies Axioms 1 to 4 and Axiom 9 he acts as if he calculates the expected utility of an act using a capacity instead of a prior as

⁵A more elaborate explanation can be found in Gilboa (2009).

in the classical Anscombe-Aumann setup. To prove Theorem 4.5 we define a capacity $v : 2^\Omega \rightarrow \mathbb{R}$ by $v(X) := \psi(\mathbb{1}_X)$ for every $X \subseteq \Omega$, where $\mathbb{1}_X$ is the vector in $[-1, 1]^\Omega$ with 1's where the respective state is included in $X \subseteq \Omega$ and 0's otherwise. We then show that computing the Choquet Integral with respect to v is equivalent to the representation obtained in Theorem 4.2. This approach is motivated by the idea that a capacity can be interpreted as a set of $n!$ priors—a relationship that is well-known in the literature. Therefore, we define the set of all acts in the same comonotonic region as $g \in \mathcal{F}$, by $\mathcal{F}_{Co}(g) := \{f \in \mathcal{F} : f \text{ is comonotonic to } g\}$. Proposition 4.8 establishes that we only need one prior for each possible comonotonic region. Moreover, there are as many comonotonic regions as there are strict ordering of the states in Ω .

Proposition 4.8. *Assume that \succeq satisfies Axioms 1 to 4 and Axiom 9. Let \succeq be represented by ψ and u as defined in Proposition 4.1, then for any act $g \in \mathcal{F}$ and all $f \in \mathcal{F}_{Co}(g)$ it holds that $U(f) = \nabla\psi(G)F$, where $F = u \circ f$ and $G = u \circ g$.*

Proof. Fix some $g \in \mathcal{F}$ and take any $f \in \mathcal{F}_{Co}(g)$. By definition of $\mathcal{F}_{Co}(g)$ the two acts f and g are comonotonic. Again, we show that $\nabla\psi(G)F = U(f)$:

$$\begin{aligned} \nabla\psi(G)F &= \psi'(G, F) \\ &= \lim_{\alpha \downarrow 0} \frac{\psi(G + \alpha F) - \psi(G)}{\alpha} \\ &= \lim_{\alpha \downarrow 0} \frac{\psi(G) + \alpha\psi(F) - \psi(G)}{\alpha} \\ &= \psi(F) \\ &= U(f), \end{aligned}$$

where the third equality follows by Comonotonic Additivity. \square

Since we have shown that Axioms 5 and 6 are weaker than Axiom 9, we can represent a decision makers' preferences also via Theorem 4.2. Note that there are $n!$ orderings of the n states in Ω . Figure 4.4.3 illustrates the resulting partition of utility vectors for the three state case. Thus, the set of acts \mathcal{F} can be partitioned into $n!$ sets such that $\mathcal{F} = \cup_{i=1}^{n!} \mathcal{F}_{Co}(g_i)$, where each $g_i \in \mathcal{F}$ represents another strict ordering of states. By Proposition 4.8, each act in a set $\mathcal{F}_{Co}(g_i)$ can be evaluated by the same prior and thus the set of priors contains at most $n!$ different priors. Since both representations are equivalent, we can simply interpret the capacity v as a set of $n!$ priors.

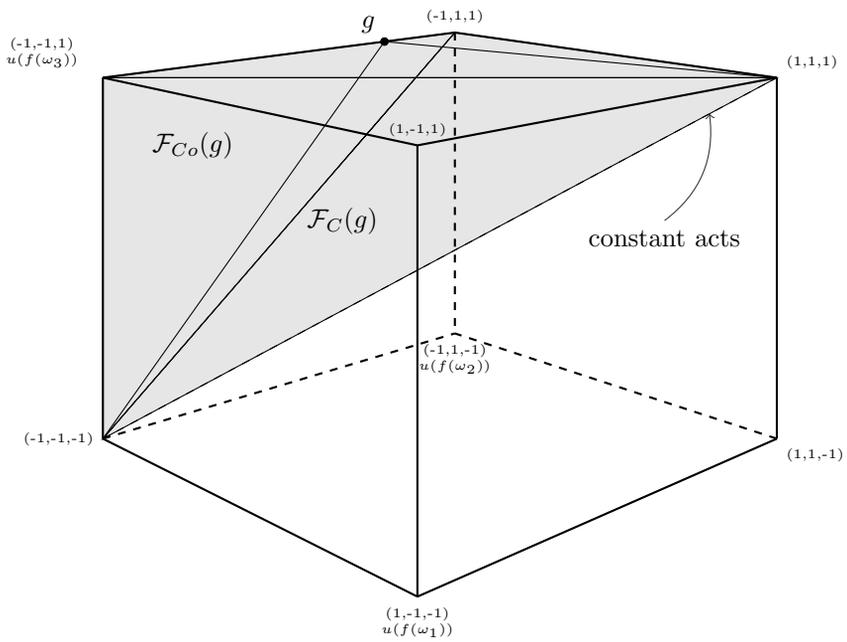


Figure 4.4.4: Example of C-independent region in a comonotonic region for three states.

Next we discuss the relationship between C-independent and comonotonic regions. First we show that for any $g \in \mathcal{F}$ we have $\mathcal{F}_C(g) \subseteq \mathcal{F}_{Co}(g)$, a fact that is also illustrated in Figure 4.4.4. Take $f \in \mathcal{F}_C(g)$ and note that f and g are comonotonic (as seen in the proof of Proposition 4.7). Thus, $f \in \mathcal{F}_{Co}(g)$ and therefore $\mathcal{F}_C(g) \subseteq \mathcal{F}_{Co}(g)$. The converse, however, is not true since not every act that is comonotonic with g is in $\mathcal{F}_C(g)$. In the two-state case, we have that $\mathcal{F}_C(g) = \mathcal{F}_{Co}(g)$ and since there are just two comonotonic regions we only require two priors. Moreover, if $n > 2$ we can take an act $g' \in \mathcal{F}$ such that g' is comonotonic to g and such that $\mathcal{F}_C(g) \cap \mathcal{F}_C(g') = \emptyset$. Therefore $g' \in \mathcal{F}_{Co}(g)$ but $g' \notin \mathcal{F}_C(g)$. However, for any $f \in \mathcal{F}_C(g')$ we have that $f \in \mathcal{F}_{Co}(g)$. Hence, we have found two sets $\mathcal{F}_C(g), \mathcal{F}_C(g') \subseteq \mathcal{F}_{Co}(g)$. This shows that there can be multiple C-independent regions in a comonotonic region.

4.4.3 Anscombe-Aumann

Anscombe and Aumann (1963) extended Von Neumann and Morgenstern's (1944) Risk Independence fully to acts. The result is the well-known Full Independence Axiom. Again, from here on we assume that \succeq satisfies Axiom 10.

Axiom 10 (Full Independence). *For any three acts $f, g, h \in \mathcal{F}$ and any $\alpha \in (0, 1)$, we have that $f \succeq g$ if and only if $\alpha f + (1 - \alpha)h \succeq \alpha g + (1 - \alpha)h$.*

Full Independence enables the standard subjective expected utility representation that is vastly used in economics. Its power is rooted in the fact that it induces Full Additivity of ψ and therefore also of the utility function U .

Proposition 4.9 (Full Additivity). *Assume that \succeq satisfies Axioms 1-4 and Axiom 10. Let \succeq be represented by ψ and u as described in Proposition 4.1, then for any two acts $f, g \in \mathcal{F}$ and $F, G \in [-1, 1]^\Omega$ such that $u \circ f = F$ and $u \circ g = G$, it holds that $\psi(F + \alpha G) = \psi(F) + \alpha\psi(G)$ for any $\alpha > 0$ with $F + \alpha G \in [-1, 1]^\Omega$.*

Given Full Additivity we can easily show that a single prior is sufficient to obtain the utility of all acts.

Proposition 4.10. *Assume that \succeq satisfies Axioms 1-4 and Axiom 10. Let \succeq be represented by ψ and u as described in Proposition 4.1, then for any act $g \in \mathcal{F}$ and all $f \in \mathcal{F}$ it holds that $U(f) = \nabla\psi(G)F$, where $F = u \circ f$ and $G = u \circ g$.*

Proof. We have to show that $U(f) = \nabla\psi(G)F$ holds for any $f, g \in \mathcal{F}$ and $\nabla\psi(G) \in \mathcal{P}_G$. We assume \succeq satisfies Axiom 10 and thereby that ψ is Fully

Additive and homogenous. Then take any $f, g \in \mathcal{F}$ and note that

$$\begin{aligned}
 \nabla\psi(G)F &= \psi'(G, F) \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi(G + \alpha F) - \psi(G)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi(G) + \alpha\psi(F) - \psi(G)}{\alpha} \\
 &= \psi(F) \\
 &= U(f),
 \end{aligned}$$

where the third equality holds by Full Additivity. □

Since Full Independence clearly implies Risk Independence as well as Independence of Certainty Equivalents, we can also represent the decision makers preferences using Theorem 4.2. Furthermore, if we combine Proposition 4.10 with Theorem 4.2 it is immediate that we only need one prior to represent the decision makers' preferences and thus we obtain the classical Anscombe-Aumann representation.

4.5 Discussion

Our approach is not the first to characterize the set of priors under weaker independence axioms. We relate our work to previous attempts and show how they connect. Next to the characterization of the set of priors, we pinpoint the exact independence assumption we need to obtain preferences that can be represented by MEU. The literature discusses even weaker assumptions on independence of acts and the implications. We show what happens in our setting if we relax C-Independence even further.

4.5.1 Characterizations of the Set of Priors

Given Axioms 1 to 6, we can characterize the set of priors by using the Gâteaux derivative of the function ψ . Ghirardato et al. (2004) use a different approach to characterize the set of priors. They define an "unambiguously preferred to" relationship, which is induced by \succeq .

Definition 4.5. Let $f, g \in \mathcal{F}$. Then f is unambiguously preferred to g , denote $f \succeq' g$, if

$$\lambda f + (1 - \lambda)h \succeq \lambda g + (1 - \lambda)h$$

for all $\lambda \in (0, 1]$ and all $h \in \mathcal{F}$.

It is clear that \succeq' is in general an incomplete relationship. They show that \succeq' can be represented in the style of Bewley (2002) using a set of priors.

Proposition 4.11. *Assume that \succeq satisfies Axioms 1-4 and Axiom 7. Let \succeq be represented by u as described in Proposition 4.1, then there exists a unique nonempty, weak compact⁶ and convex set \mathcal{C} of priors on Ω such that for all $f, g \in \mathcal{F}$,*

$$f \succeq' g \iff \int_{\Omega} u(f)dp \geq \int_{\Omega} u(g)dp \text{ for all } p \in \mathcal{C}.$$

Furthermore, they show that the set of priors \mathcal{C} is equal to the Clarke differential of ψ at 0, which is a generalized derivative for functions that are not differentiable. An important difference to a derivative is that the Clarke differential is set valued. Since the Gâteaux derivative of ψ is bound if it exists, ψ is locally Lipschitz and therefore by Ghirardato et al.'s (2004) Lemma A.7 we have that $\partial\psi(0) = \mathcal{P}$, where $\partial\psi$ is the Clarke differential of ψ and \mathcal{P} is as defined in Section 4.4.1. Hence, both representations identify the same set of relevant priors.

4.5.2 Further Weakening of the C-Independence Axiom

Axiom 7, C-Independence, entails two features: (i) invariance to rescaling of utilities (by choosing $\langle p \rangle$ with $u(p) = 0$) and (ii) invariance to translations of utility profiles. Recall that a preference relation \succeq satisfying Axioms 1 to 5 allows a representation as in Proposition 4.1. Nevertheless, by requiring (i) we can only show that ψ is positively homogeneous, which is insufficient to show that $\nabla\psi(F)$ is a prior for all $F \in [-1, 1]^{\Omega}$. The same problem prevails when we completely dispose of any independence assumption on acts. Thus C-Additivity, which is ensured by (ii) as we will show below, seems to be necessary to obtain a simple prior representation. Maccheroni et al. (2006) present the following axiom that guarantees (ii) but not necessarily (i).

⁶Compact with respect to the weak topology.

Axiom 11 (Weak Certainty Independence). *If $f, g \in \mathcal{F}$, $p, q \in \Delta(X)$, and $\lambda \in (0, 1)$ then $\lambda f + (1 - \lambda)\langle p \rangle \succeq \lambda g + (1 - \lambda)\langle p \rangle$ implies $\lambda f + (1 - \lambda)\langle q \rangle \succeq \lambda g + (1 - \lambda)\langle q \rangle$.*

Thus, given that an act mixed with a constant act is preferred to another act mixed with the same constant act, the first act is always preferred to the second when mixed in the same proportion with another constant act. For a preference relation \succeq that satisfies Axioms 1 to 5 Proposition 4.1 holds and Axiom 11 induces C-Additivity of the function ψ but not homogeneity.

Proposition 4.12 (C-Additive). *Assume that \succeq satisfies Axioms 1-5 and Axiom 11. Let \succeq be represented by ψ and u as described in Proposition 4.1, then $\psi(F + \alpha \mathbb{1}_\Omega) = \psi(F) + \alpha \psi(\mathbb{1}_\Omega)$ for $\alpha > 0$ where $F + \alpha \mathbb{1}_\Omega \in [-1, 1]^\Omega$.*

The lack of homogeneity of ψ prevents an equivalent of Lemma 4.2, which establishes that the $\psi(F)$ can be computed by $\psi'(F, F)$ for $F \in [-1, 1]^\Omega$. However, we can use the following generalized Mean Value Theorem to show that $\psi(F) = \psi'(F - sF, F)$ for some $s \in (0, 1)$ by simply letting $G := -F$ for $F \in \text{int}([-1, 1]^\Omega)$. For some F that is not in the interior of $[-1, 1]^\Omega$, we can approximate the utility by taking a sequence in $\text{int}([-1, 1]^\Omega)$ that converges to F , which is possible by Continuity.

Theorem 4.6 (Generalized Mean Value Theorem). *Let K be a nonempty open subset of $[-1, 1]^\Omega$ and ψ be Gâteaux differentiable with Gâteaux derivative $\psi'(F, G)$ at $F \in [-1, 1]^\Omega$ in the direction of $G \in [-1, 1]^\Omega$. Then for any $F \in K$ and $F + G \in K$, there exists some $s \in (0, 1)$ such that*

$$\psi(F + G) - \psi(F) = \psi'(F + sG, G).$$

Proof. For the proof consult Ansari et al. (2013). □

By Lemma 4.1 we know that $\psi'(F - sF, F)$ can be represented as $\nabla\psi(F - sF)F$, if the Gâteaux derivative of ψ exists at $F - sF$. Again, it turns out that $\nabla\psi(F - sF)$ is a prior.

Proposition 4.13. *Assume that \succeq satisfies Axioms 1-5 and Axiom 11. Let \succeq be represented by ψ and u as described in Proposition 4.1, then if the Gâteaux derivative of ψ at $F \in [-1, 1]^\Omega$ exists, the gradient $\nabla\psi(F)$ of ψ has the following properties:*

- $\nabla\psi(F)(\omega) \geq 0$ for all $\omega \in \Omega$, and

- $\sum_{\omega \in \Omega} \nabla \psi(F)(\omega) = 1$.

The proof of Proposition 4.13 is equivalent to that of Proposition 4.3 since the proof only uses Monotonicity, homogeneity of ψ' in the second component and C-Additivity. The following representation theorem is thus immediate.

Theorem 4.7. *Assume that \succeq satisfies Axioms 1-5 and Axiom 11. Let \succeq be represented by ψ and u as described in Proposition 4.1, then if $F \in \text{int}([-1, 1]^\Omega)$, the Gâteaux derivative of ψ at F exists, and $F = u \circ f$ for some $f \in \mathcal{F}$, then*

$$U(f) = \nabla \psi(\hat{F})F,$$

where $\nabla \psi(\hat{F})$ is the gradient at $\hat{F} := F - sF$ for some $s \in (0, 1)$. Moreover, $\nabla \psi(\hat{F})$ is a probability measure.

Hence, the decisions maker still evaluates acts using a prior, however, we cannot say which prior exactly since s is possibly different for every act $f \in \mathcal{F}$.

2.A Appendix - Proofs Section 4.1

From now on let \succeq be the agent's preference relation over \mathcal{F} . Then let \succeq^* be the preference relation over $\Delta(X)$ induced by \succeq such that for any $p, q \in \Delta(X)$ we have $p \succeq^* q$ if and only if $\langle p \rangle \succeq \langle q \rangle$.

Proposition 4.14. *Assume that \succeq satisfies Axioms 1-5, then the preference relation \succeq^* induced by \succeq satisfies Weak Order, Continuity, and Risk Independence with respect to lotteries in $\Delta(X)$.*

Proof. The properties of \succeq^* follow directly from the definition. □

Proposition 4.15. *Assume that \succeq satisfies Axioms 1-5, then there exist $p_*, p^* \in \Delta(X)$ such that for all acts $f \in \mathcal{F}$ we have that $\langle p^* \rangle \succeq f \succeq \langle p_* \rangle$.*

Proof. Let $x^*, x_* \in X$ be such that $\langle x^* \rangle \succeq \langle x \rangle \succeq \langle x_* \rangle$ for all $x \in X$, where $\langle x \rangle$ is the constant act with $\langle x \rangle(\omega)(x) = 1$ for all $\omega \in \Omega$. Then by Axiom 5 it follows that $\langle x^* \rangle \succeq \langle p \rangle \succeq \langle x_* \rangle$ for all $p \in \Delta(X)$. By Axiom 2, $\langle x^* \rangle \succeq f \succeq \langle x_* \rangle$ for all $f \in \mathcal{F}$. Set $\langle p^* \rangle := \langle x^* \rangle$ and $\langle p_* \rangle := \langle x_* \rangle$. Thus, we have for all acts $f \in \mathcal{F}$ that $\langle p^* \rangle \succeq f \succeq \langle p_* \rangle$. □

Proposition 4.16. *Assume that \succeq satisfies Axioms 1-5, then there exists an affine map u on $\Delta(X)$ that represents \succeq^* .*

Proof. Since \succeq^* satisfies Weak Order, Continuity, and Risk Independence it follows from the Von Neumann-Morgenstern Expected Utility Theorem (Von Neumann and Morgenstern, 1944) that there exists an affine map u on $\Delta(X)$ that represents \succeq^* . □

Proposition 4.17. *Assume that \succeq satisfies Axioms 1-5, then for any $f \in \mathcal{F}$, there is a unique $\alpha_f \in [0, 1]$ such that $f \sim \alpha_f \langle p^* \rangle + (1 - \alpha_f) \langle p_* \rangle$.*

Proof. First we have to show that for any act $f \in \mathcal{F}$ there exists a mixture of constant acts $p^*, p_* \in \Delta(X)$ such that $f \sim \alpha \langle p^* \rangle + (1 - \alpha) \langle p_* \rangle$ for some $\alpha \in [0, 1]$. If $f \sim \langle p^* \rangle$ we simply take $\alpha = 1$ and if $f \sim \langle p_* \rangle$ we take $\alpha = 0$. In the case in which $\langle p^* \rangle \succ f \succ \langle p_* \rangle$, Axiom 3 guarantees that there exists a $\alpha \in (0, 1)$ such that $f \sim \alpha \langle p^* \rangle + (1 - \alpha) \langle p_* \rangle$.

Secondly, we have to show that the α for which $f \sim \alpha \langle p^* \rangle + (1 - \alpha) \langle p_* \rangle$ holds is unique. Take the acts $p^*, p_* \in \Delta(X)$. Then, $\langle p^* \rangle \succ \langle p_* \rangle$ due to Axiom 4.

Subsequently, apply Axiom 5 to obtain for any $\lambda \in (0, 1)$

$$\lambda \langle p^* \rangle + (1 - \lambda) \langle p_* \rangle \succ \lambda \langle p_* \rangle + (1 - \lambda) \langle p^* \rangle,$$

where the left-hand side can simply be expressed as $\langle p^* \rangle$. Then we apply Axiom 5 again and obtain for any $\gamma \in (0, 1)$

$$\gamma \langle p^* \rangle + (1 - \gamma) \langle p_* \rangle \succ (\gamma(1 - \lambda)) \langle p^* \rangle + (1 - \gamma(1 - \lambda)) \langle p_* \rangle. \quad (2.A.1)$$

Now we take any $\alpha, \beta \in (0, 1)$ and assume w.l.o.g. $\alpha > \beta$. Then take $\gamma = \alpha$ and $\lambda = (\gamma - \beta)/\gamma$. Then $\beta = \gamma(1 - \lambda)$ and equation (2.A.1) implies that $\alpha \langle p^* \rangle + (1 - \alpha) \langle p_* \rangle \succ \beta \langle p^* \rangle + (1 - \beta) \langle p_* \rangle$. Since this is true for any $\alpha > \beta$ we can conclude that α is unique. \square

Proposition 4.18. *Assume that \succeq satisfies Axioms 1-5, then there exists a unique $U : \mathcal{F} \rightarrow \mathbb{R}$ that represents \succeq and satisfies*

$$U(\langle p \rangle) = u(p) \quad (2.A.2)$$

for all $p \in \Delta(X)$.

Proof. Define U on \mathcal{F} by $U(f) := u(\alpha_f p^* + (1 - \alpha_f) p_*)$, where α_f is found as in Proposition 4.17. It's clear that U represents \succeq and satisfies (2.A.2). Moreover, if V was another such function, we would have for any $f \in \mathcal{F}$,

$$U(f) = U(\langle p_f \rangle) = u(p_f) = V(\langle p_f \rangle) = V(f),$$

where $p_f = \alpha_f p^* + (1 - \alpha_f) p_*$. \square

From now on we assume $u(p^*) = 1$ and $u(p_*) = -1$.

Proposition 4.19. *Assume that \succeq satisfies Axioms 1-5, then $\{u \circ f : f \in \mathcal{F}\} = [-1, 1]^\Omega$.*

Proof. For the " \subseteq "-direction remember that $u : \Delta(X) \rightarrow [-1, 1]$. Since the composition is a state-wise operation, we have $u \circ f \in [-1, 1]^\Omega$ for all $f \in \mathcal{F}$.

Now consider the " \supseteq "-direction. Take a $F \in [-1, 1]^\Omega$. Then since Axiom 3 holds for u by the Mean Value Theorem there exists a $p_\omega \in \Delta(X)$ such that $u(p_\omega) = F(\omega)$ for every $\omega \in \Omega$. Let $g \in \mathcal{F}$ be such that $g(\omega) := p_\omega$ and therefore we have $u \circ g = F$. \square

Definition 4.6. Let the function $\psi : [-1, 1]^\Omega \rightarrow \mathbb{R}$ be such that $\psi(u \circ f) := U(f)$ for all $f \in \mathcal{F}$.

Such a function ψ must exist because Axiom 2 implies that $U(f) = U(g)$ whenever $u \circ f = u \circ g$. Since U represents \succeq and we defined $\psi(u \circ f) = U(f)$ for all $f \in \mathcal{F}$, the function ψ is clearly monotonic and continuous.

Proof Proposition 4.1. By Proposition 4.16 there exists an affine map u on $\Delta(X)$ that represents \succeq^* . Take $p^*, p_* \in \Delta(X)$ as in Proposition 4.15 and define $u(p^*) = 1$ and $u(p_*) = -1$. Then according to Proposition 4.18 there exists a utility function U on \mathcal{F} that represents \succeq and is such that $U(\langle p \rangle) = u(p)$ for all $p \in \Delta(X)$. Moreover, we know that there exists a function $\psi : [-1, 1]^\Omega \rightarrow \mathbb{R}$ such that $\psi(u \circ f) := U(f)$ for all $f \in \mathcal{F}$. Since ψ also represents \succeq it is continuous and monotonic. Finally, it follows directly from Proposition 4.17 that for all $f \in \mathcal{F}$ there is a $p \in \Delta(X)$ such that $f \sim \langle p \rangle$. \square

2.B Appendix - Proofs Section 4.2

Proof Lemma 4.1. Take $F \in [-1, 1]^\Omega$ where the Gâteaux derivative of ψ exists and any $G \in [-1, 1]^\Omega$. Then,

$$\begin{aligned} \nabla\psi(F)G &= \sum_{\omega \in \Omega} \psi'(F, \mathbb{1}_\omega)G(\omega) \\ &= \psi'(F, \sum_{\omega \in \Omega} G(\omega)\mathbb{1}_\omega) \\ &= \psi'(F, G), \end{aligned}$$

where the second equality holds by homogeneity of ψ' in the second component. \square

2.C Appendix - Proofs Section 4.3

Proof Proposition 4.2. For the first property take any $f \in \mathcal{F}$ and $F \in [-1, 1]^\Omega$ such that $u \circ f = F$. Let $g := \lambda f + (1 - \lambda)\langle p \rangle$ with $u(p) = 0$ and $\lambda \in (0, 1]$. Note that

$$\begin{aligned} u \circ g &= u \circ (\lambda f + (1 - \lambda)\langle p \rangle) \\ &= \lambda(u \circ f) + (1 - \lambda)u \circ \langle p \rangle \\ &= \lambda(u \circ f), \end{aligned}$$

where the second equality holds by affinity of u . Applying ψ yields $\psi(u \circ g) = \psi(\lambda(u \circ f)) = \psi(\lambda F)$. By Proposition 4.1 we can find a $q \in \Delta(X)$ such that $f \sim \langle q \rangle$. Thus, Axiom 6 implies $g \sim \lambda\langle q \rangle + (1 - \lambda)\langle p \rangle$. The utility of g can then be expressed as

$$\begin{aligned} U(g) &= u(\lambda q + (1 - \lambda)p) \\ &= \lambda u(q) \\ &= \lambda U(\langle q \rangle) \\ &= \lambda U(f), \end{aligned}$$

where again the second equality follows by affinity of u . Now we can show that $U(g) = \lambda\psi(F)$ by noting that

$$U(g) = \lambda U(f) = \lambda\psi(u \circ f) = \lambda\psi(F),$$

and thus that

$$\psi(\lambda F) = \psi(\lambda(u \circ f)) = \psi(u \circ g) = U(g) = \lambda\psi(F),$$

for all $\lambda \in (0, 1]$. Finally, we show that $\psi(\lambda F) = \lambda\psi(F)$ holds for all $\lambda > 1$ as well. Take any $F \in [-1, 1]^\Omega$ and $\lambda > 1$ such that $\lambda F \in [-1, 1]^\Omega$. Then it follows from the above that $\psi(F) = \psi(\frac{1}{\lambda}\lambda F) = \frac{1}{\lambda}\psi(\lambda F)$ and hence $\lambda\psi(F) = \psi(\lambda F)$ for all $\lambda > 1$.

To prove the second property take any $f \in \mathcal{F}$ and $F \in [-1, 1]^\Omega$ such that $u \circ f = F$. Furthermore, take $\alpha \in [0, 1]$ and $p \in \Delta(X)$ such that $u(p) = \alpha$. By Proposition 4.1 we can find $q \in \Delta(X)$ such that $f \sim \langle q \rangle$. Therefore, Axiom 6 implies $\frac{1}{2}f + \frac{1}{2}\langle p \rangle \sim \frac{1}{2}\langle q \rangle + \frac{1}{2}\langle p \rangle$. It follows immediately that $U(\frac{1}{2}f + \frac{1}{2}\langle p \rangle) =$

$U(\frac{1}{2}\langle q \rangle + \frac{1}{2}\langle p \rangle)$. We now use positive homogeneity of ψ and affinity of u to show

$$\begin{aligned} 2U\left(\frac{1}{2}f + \frac{1}{2}\langle p \rangle\right) &= 2\psi\left(u \circ \left(\frac{1}{2}f + \frac{1}{2}\langle p \rangle\right)\right) \\ &= \psi(u \circ f + u \circ \langle p \rangle) \\ &= \psi(F + \alpha \mathbb{1}_\Omega). \end{aligned}$$

Furthermore, we can use u 's affinity and the definitions of U and ψ to show

$$\begin{aligned} 2U\left(\frac{1}{2}\langle q \rangle + \frac{1}{2}\langle p \rangle\right) &= 2u\left(\frac{1}{2}q + \frac{1}{2}p\right) \\ &= u(q) + \alpha \\ &= U(\langle q \rangle) + \alpha \\ &= U(f) + \alpha \\ &= \psi(F) + \alpha. \end{aligned}$$

Combining our two findings and using the fact that $U(\frac{1}{2}f + \frac{1}{2}\langle p \rangle) = U(\frac{1}{2}\langle q \rangle + \frac{1}{2}\langle p \rangle)$ we obtain $\psi(F + \alpha \mathbb{1}_\Omega) = \psi(F) + \alpha\psi(\mathbb{1}_\Omega)$ since $\psi(\mathbb{1}_\Omega) = 1$. Finally, we show that $\psi(F + \alpha \mathbb{1}_\Omega) = \psi(F) + \alpha$ holds for all $\alpha > 1$ as well. Take $F \in [-1, 1]^\Omega$ and $\alpha > 1$ such that $F + \alpha \mathbb{1}_\Omega \in [-1, 1]^\Omega$. Note that $\alpha \leq 2$. Let $G := F + \mathbb{1}_\Omega$, then we have

$$\begin{aligned} \psi(F + \alpha \mathbb{1}_\Omega) &= \psi(F + \mathbb{1}_\Omega - \mathbb{1}_\Omega + \alpha \mathbb{1}_\Omega) \\ &= \psi(G + (\alpha - 1) \mathbb{1}_\Omega) \\ &= \psi(G) + (\alpha - 1) \\ &= \psi(F + \mathbb{1}_\Omega) + (\alpha - 1) \\ &= \psi(F) + \alpha, \end{aligned}$$

where the third and fifth equality hold by the result obtained above. \square

Proof Lemma 4.2. We start by considering the gradient of ψ at any $F \in [-1, 1]^\Omega$ where the Gâteaux derivative exists. We can use the definition ψ' and the homo-

ogeneity of ψ to show that

$$\begin{aligned}
 \psi'(F, F) &= \lim_{\alpha \downarrow 0} \frac{\psi(F + \alpha F) - \psi(F)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi((1 + \alpha)F) - \psi(F)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{(1 + \alpha)\psi(F) - \psi(F)}{\alpha} \\
 &= \psi(F).
 \end{aligned}$$

□

Proof Proposition 4.3. Consider the first property and therefore the gradient of ψ at $F \in [-1, 1]^\Omega$ at state $\omega \in \Omega$:

$$\begin{aligned}
 \nabla\psi(F)(\omega) &= \psi'(F, \mathbb{1}_\omega) \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi(F + \alpha \mathbb{1}_\omega) - \psi(F)}{\alpha}.
 \end{aligned}$$

By Axiom 2 we have $\psi(F + \alpha \mathbb{1}_\omega) \geq \psi(F)$ for all $\alpha \geq 0$ and thus the limit needs to be larger or equal to zero. It follows that $\nabla\psi(F)(\omega) \geq 0$.

For the second property, consider the sum

$$\begin{aligned}
 \sum_{\omega \in \Omega} \nabla\psi(F)(\omega) &= \sum_{\omega \in \Omega} \psi'(F, \mathbb{1}_\omega) \\
 &= \psi'(F, \sum_{\omega \in \Omega} \mathbb{1}_\omega) \\
 &= \psi'(F, \mathbb{1}_\Omega) \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi(F + \alpha \mathbb{1}_\Omega) - \psi(F)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{\psi(F) + \alpha\psi(\mathbb{1}_\Omega) - \psi(F)}{\alpha} \\
 &= \lim_{\alpha \downarrow 0} \frac{\alpha\psi(\mathbb{1}_\Omega)}{\alpha} \\
 &= \psi(\mathbb{1}_\Omega) = 1.
 \end{aligned}$$

The second equality follows by homogeneity of ψ' in the second component and the fifth by C-Additivity of ψ . □

Proof Proposition 4.4. Take any act $f \in \mathcal{F}$ and $F \in [-1, 1]^\Omega$ such that $u \circ f = F$, where the derivative of ψ at F does not exist. Since ψ is continuous and the derivative of ψ exists almost everywhere, we can approximate the utility of f by taking a sequence $(F_n)_{n \in \mathbb{N}}$ such that the Gâteaux derivative exists at all F_n for all n and $F_n \in [-1, 1]^\Omega$. Furthermore, we require $\lim_{n \rightarrow \infty} F_n = F$. Then, by continuity of ψ , $\psi(F) = \lim_{n \rightarrow \infty} \psi(F_n) = \lim_{n \rightarrow \infty} \nabla \psi(F_n) F_n$, where the last equality follows from Theorem 4.2. \square

Proof of Theorem 4.3. The "only if" direction follows from Theorem 4.2, and Proposition 4.1. Thus we only have to prove the "if" direction. Axiom 1 (Weak Order) is implied by the existence of a utility function U .

For Axiom 2 (Monotonicity) take $f, g \in \mathcal{F}$ such that $\langle f(\omega) \rangle \succeq \langle g(\omega) \rangle$ for all $\omega \in \Omega$ and let $F = u \circ f$ and $G = u \circ g$. Then $F(\omega) \geq G(\omega)$ for all $\omega \in \Omega$. By monotonicity of ψ we have $\psi(F) \geq \psi(G)$ and since $\psi(F) = U(f)$ for all $f \in \mathcal{F}$ it follows that $f \succeq g$.

To show that \succeq satisfies Axiom 3, take $f, g, h \in \mathcal{F}$ such that $f \succ g \succ h$. We have to show that there is an $\alpha \in (0, 1]$ such that $U(g) = U(\alpha f + (1 - \alpha)g)$. Let $F = u \circ f, G = u \circ g, H = u \circ h$. Then, we must show that there is an $\alpha \in (0, 1]$ such that $\psi(G) = \psi(\alpha F + (1 - \alpha)H)$. This, however, follows from continuity of ψ and the fact that $\psi(F) > \psi(G)$ as well as $\psi(H) < \psi(G)$.

Axiom 4 simply follows from the assumption that there exist $p_*, p^* \in \Delta(X)$ such that $u(p^*) = 1$ and $u(p_*) = -1$.

For Axiom 5 take $x, y, z \in \Delta(X)$ such that $\langle x \rangle \succ \langle y \rangle$. Moreover, let $u(x) = \alpha$, $u(y) = \beta$, and $u(z) = \gamma$. Take some $\lambda \in (0, 1]$ and note that

$$\begin{aligned}
 U(\lambda\langle x \rangle + (1 - \lambda)\langle z \rangle) &= \psi(u \circ (\lambda\langle x \rangle + (1 - \lambda)\langle z \rangle)) \\
 &= \psi(\lambda u \circ \langle x \rangle + (1 - \lambda)u \circ \langle z \rangle) \\
 &= \psi(\lambda\alpha\mathbb{1}_\Omega + (1 - \lambda)\gamma\mathbb{1}_\Omega) \\
 &= \lambda\psi(\alpha\mathbb{1}_\Omega) + (1 - \lambda)\psi(\gamma\mathbb{1}_\Omega) \\
 &= \lambda U(\langle x \rangle) + (1 - \lambda)U(\langle z \rangle) \\
 &> \lambda U(\langle y \rangle) + (1 - \lambda)U(\langle z \rangle) \\
 &= \lambda\psi(\beta\mathbb{1}_\Omega) + (1 - \lambda)\psi(\gamma\mathbb{1}_\Omega) \\
 &= \psi(\lambda\beta\mathbb{1}_\Omega + (1 - \lambda)\gamma\mathbb{1}_\Omega) \\
 &= U(\lambda\langle y \rangle + (1 - \lambda)\langle z \rangle),
 \end{aligned}$$

where the second equality holds by affinity of u and the fourth and seventh equality hold by positive homogeneity and C-Additivity of ψ , implying $\lambda\langle x \rangle + (1 - \lambda)\langle z \rangle \succ \lambda\langle y \rangle + (1 - \lambda)\langle z \rangle$.

Finally, for Axiom 6 we have to show that if $f \sim \langle p \rangle$, where $f \in \mathcal{F}$ and $p \in \Delta(X)$, then $U(\lambda f + (1 - \lambda)\langle q \rangle) = U(\lambda\langle p \rangle + (1 - \lambda)\langle q \rangle)$ for all $q \in \Delta(X)$ and $\lambda \in [0, 1)$. Take $f \in \mathcal{F}$ and $p \in \Delta(X)$ with $f \sim \langle p \rangle$ and $u(p) = \alpha$, and some $q \in \Delta(X)$ such that $u(q) = \beta$. Let $F = u \circ f$, $\lambda \in [0, 1)$ and note that

$$\begin{aligned}
 U(\lambda f + (1 - \lambda)\langle q \rangle) &= \psi(\lambda F + (1 - \lambda)\beta\mathbb{1}_\Omega) \\
 &= \lambda\psi(F) + \psi((1 - \lambda)\beta\mathbb{1}_\Omega) \\
 &= \lambda U(f) + U((1 - \lambda)\langle q \rangle) \\
 &= \lambda U(\langle p \rangle) + U((1 - \lambda)\langle q \rangle) \\
 &= \lambda\psi(\alpha\mathbb{1}_\Omega) + \psi((1 - \lambda)\beta\mathbb{1}_\Omega) \\
 &= \psi(\lambda\alpha\mathbb{1}_\Omega + (1 - \lambda)\beta\mathbb{1}_\Omega) \\
 &= U(\lambda\langle p \rangle + (1 - \lambda)\langle q \rangle),
 \end{aligned}$$

where the second and the sixth equality hold by positive homogeneity and C-Additivity of ψ . □

2.D Appendix - Proofs Section 4.4

Proof of Proposition 4.5. First of all, take any $f, g \in \mathcal{F}$ such that $f \sim g$. Then by Axiom 8 we have that $\frac{1}{2}f + \frac{1}{2}g \succeq f$. Let $F := u \circ f$ and $G := u \circ g$, then it follows that

$$\begin{aligned} \psi\left(\frac{1}{2}F + \frac{1}{2}G\right) &= \psi\left(u \circ \left(\frac{1}{2}f + \frac{1}{2}g\right)\right) \\ &= U\left(\frac{1}{2}f + \frac{1}{2}g\right) \\ &\geq U(f) \\ &= \psi(F) \\ &= \frac{1}{2}\psi(F) + \frac{1}{2}\psi(G). \end{aligned}$$

Then by positive homogeneity of ψ we have that $\psi(F + G) \geq \psi(F) + \psi(G)$. Now consider $f, g \in \mathcal{F}$ such that $f \succ g$. Let $\alpha := \psi(F) - \psi(G)$ and define $H := G + \alpha \mathbb{1}_\Omega$. By C-Additivity and the fact that $\psi(\mathbb{1}_\Omega) = 1$ we have $\psi(H) = \psi(G) + \alpha = \psi(F)$. Thus by our previous result it holds that

$$\psi(F + G + \alpha \mathbb{1}_\Omega) = \psi(F + H) \geq \psi(F) + \psi(H) = \psi(F) + \psi(G) + \alpha.$$

Then by C-Additivity and homogeneity of ψ we have $\psi(F + G + \alpha \mathbb{1}_\Omega) = \psi(F + G) + \psi(\alpha \mathbb{1}_\Omega) = \psi(F + G) + \alpha$. Therefore, it holds that $\psi(F + G) \geq \psi(F) + \psi(G)$. \square

Proposition 4.20. *Let \succeq be a preference relation on \mathcal{F} satisfying Axioms 1 to 5, and Axiom 9 then it holds that for all pairwise comonotonic acts $f, g, h \in \mathcal{F}$ with $f \succeq g$ that $\alpha f + (1 - \alpha)h \succeq \alpha g + (1 - \alpha)h$ for all $\alpha \in (0, 1)$.*

Proof. First of all, note that Axiom 9 implies Axioms 5 and 6 and thus Theorem 4.2 holds. Then, if there is no $g \in \mathcal{F}$ with $g \succ f$, take any $f, g, h \in \mathcal{F}$ such that f, g, h are pairwise comonotonic and $f \sim g$. Note that both f and g must be constant acts and therefore $u \circ f = u \circ g$. This holds because f must be the act that gives 1 in every state, otherwise there would be an $h \in \mathcal{F}$ such that $h \succ f$. Then either g is a constant act or we take g 's certainty equivalent and simply call it g . Since $f \sim g$ the constant acts must be equal. Hence, we must have that $u \circ (\lambda f + (1 - \lambda)h) = u \circ (\lambda g + (1 - \lambda)h)$ for $\lambda \in (0, 1)$. This implies $\psi(u \circ (\lambda f + (1 - \lambda)h)) = \psi(u \circ (\lambda g + (1 - \lambda)h))$, which by affinity of u is

equivalent to $\psi(\lambda u \circ f + (1 - \lambda)u \circ h) = \psi(\lambda u \circ g + (1 - \lambda)u \circ h)$. Hence, we have that $U(\lambda f + (1 - \lambda)h) = U(\lambda g + (1 - \lambda)h)$ and therefore $\lambda f + (1 - \lambda)h \sim \lambda g + (1 - \lambda)h$. In the case where there is no $g \in \mathcal{F}$ such that $f \succ g$ then we can use the same argument.

Now consider the case where $f \in \mathcal{F}$ is in the interior of the set of acts and take any $f, g, h \in \mathcal{F}$ such that f, g and h are pairwise comonotonic and $f \sim g$. Then take two sequences such that $f_n^+ \rightarrow f$ and $f_n^- \rightarrow f$ as $n \rightarrow \infty$, such that f_n^+ and f_n^- are comonotonic to f . Moreover, we require that $f_n^+ \succ f$ and $f \succ f_n^-$ for all $n \in \mathbb{N}$. Thus, it must hold that $f_n^+ \succ g$ and therefore by Comonotonic Independence that $\lambda f_n^+ + (1 - \lambda)h \succ \lambda g + (1 - \lambda)h$ for any $\lambda \in (0, 1)$. Furthermore, it must hold that $g \succ f_n^-$ and therefore by Comonotonic Independence that $\lambda g + (1 - \lambda)h \succ \lambda f_n^- + (1 - \lambda)h$ for any $\lambda \in (0, 1)$. Hence, as $n \rightarrow \infty$ it must hold by Continuity that $f \sim g$ implies $\lambda f + (1 - \lambda)h \sim \lambda g + (1 - \lambda)h$ for any $\lambda \in (0, 1)$. This gives us the symmetric part of Comonotonic Independence and since the asymmetric part is already given in Axiom 9 it concludes the proof. \square

Lemma 4.3. *Assume that \succeq satisfies Axioms 1-6. Let \succeq be represented by ψ and u as described in Proposition 4.1, then take any $f, g \in \mathcal{F}$ and $p \in \Delta(X)$ such that $u(p) = 0$. Let $h := \alpha g + (1 - \alpha)\langle p \rangle$ for some $\alpha \in (0, 1)$ and $F, G, H \in [-1, 1]^\Omega$ such that $u \circ f = F$, $u \circ g = G$, and $u \circ h = H$. If it holds that $\frac{1}{2}f + \frac{1}{2}h \sim \frac{1}{2}f + \frac{1}{2}\langle q \rangle$, where $h \sim \langle q \rangle$ for some $q \in \Delta(X)$, then we have that*

$$\psi(F + \alpha G) = \psi(F) + \alpha\psi(G).$$

Proof. Note that $2U(\frac{1}{2}f + \frac{1}{2}h) = 2U(\frac{1}{2}f + \frac{1}{2}\langle q \rangle)$. First we take on the left-hand side and show that $2U(\frac{1}{2}f + \frac{1}{2}h) = \psi(F + \alpha G)$ by using the affinity of u and positive homogeneity of ψ :

$$\begin{aligned} 2U\left(\frac{1}{2}f + \frac{1}{2}h\right) &= 2\psi\left(u \circ \left(\frac{1}{2}f + \frac{1}{2}h\right)\right) \\ &= \psi(u \circ f + u \circ h) \\ &= \psi(u \circ f + (u \circ (\alpha g + (1 - \alpha)\langle p \rangle))) \\ &= \psi(u \circ f + \alpha u \circ g) \\ &= \psi(F + \alpha G). \end{aligned}$$

For the right-hand side we perform a similar exercise:

$$\begin{aligned}
 2U\left(\frac{1}{2}f + \frac{1}{2}\langle q \rangle\right) &= 2\psi\left(\frac{1}{2}F + \frac{1}{2}u \circ \langle q \rangle\right) \\
 &= \psi(F + u \circ \langle q \rangle) \\
 &= \psi(F) + \psi(u \circ \langle q \rangle) \\
 &= \psi(F) + \psi(u \circ (\alpha g + (1 - \alpha)\langle p \rangle)) \\
 &= \psi(F) + \psi(\alpha u \circ g + (1 - \alpha)u \circ \langle p \rangle) \\
 &= \psi(F) + \alpha\psi(u \circ g) \\
 &= \psi(F) + \alpha\psi(G).
 \end{aligned}$$

To obtain the second equality we use positive homogeneity of ψ , for the third we use C-Additivity of ψ , for the fifth we use affinity of u , and for the sixth equality we use positive homogeneity of ψ . Thus we showed that $\psi(F + \alpha G) = \psi(F) + \alpha\psi(G)$. \square

Proof of Proposition 4.7. Define $h := \alpha g + (1 - \alpha)\langle p \rangle$ for any $\alpha \in (0, 1)$ where $p \in \Delta(X)$ is such that $u(p) = 0$. First we show that f and h are comonotonic. Take $H \in [-1, 1]^\Omega$ such that $u \circ h = H$. Then we have to show that $(F(\omega) - F(\omega'))(H(\omega) - H(\omega')) \geq 0$ for $\omega, \omega' \in \Omega$, a condition which is equivalent to the definition of Comonotonicity. Note that for any two states $\omega, \omega' \in \Omega$ we have that

$$\begin{aligned}
 H(\omega) - H(\omega') &= (\alpha G(\omega) + (1 - \alpha)u(p)) - (\alpha G(\omega') + (1 - \alpha)u(p)) \\
 &= \alpha(G(\omega) - G(\omega')).
 \end{aligned}$$

Since $\alpha > 0$ it does not change the sign of $G(\omega) - G(\omega')$ and thus h is comonotonic to g and therefore f to h . Furthermore, take $q \in \Delta(X)$ such that $\langle q \rangle \sim h$. By Proposition 4.20 we obtain $\frac{1}{2}f + \frac{1}{2}h \sim \frac{1}{2}f + \frac{1}{2}\langle q \rangle$. Now we can use Lemma 4.3 to show that $\psi(F + \alpha G) = \psi(F) + \alpha\psi(G)$ for any $\alpha \in (0, 1)$. Then for $\alpha \geq 1$ and any $F, \alpha G \in [-1, 1]^\Omega$ by positive homogeneity of ψ and the result above we have

that

$$\begin{aligned}
 \psi(F + \alpha G) &= \psi\left(\alpha\left(\frac{1}{\alpha}F + G\right)\right) \\
 &= \alpha\psi\left(\frac{1}{\alpha}F + G\right) \\
 &= \alpha\psi\left(\frac{1}{\alpha}F\right) + \alpha\psi(G) \\
 &= \psi(F) + \alpha\psi(G).
 \end{aligned}$$

Thus $\psi(F + \alpha G) = \psi(F) + \alpha\psi(G)$ holds for any $\alpha > 0$. \square

Proof of Theorem 4.5. Define the capacity v by $v(X) := \psi(\mathbb{1}_X)$ for any $X \subseteq \Omega$, and $v(\emptyset) = 0$. Remember that $\mathbb{1}_X$ is an $1 \times n$ vector with 1's where the respective state is included in X . Our task is to show that v is the unique capacity such that (4.4.2) holds. For that matter, take any act $f \in \mathcal{F}$ and a $F \in [-1, 1]^\Omega$ such that $u \circ f = F$. Recall that σ is a bijective function that orders the states of F from best to worst and define $O(i) := \{\sigma(1), \dots, \sigma(i)\}$ as the set of ordered states of F up to rank i .

First of all, note that (4.4.1) can be rewritten as

$$\int_{\Omega} F dv = \sum_{i=1}^n F(\sigma(i)) (v(O(i)) - v(O(i-1))). \quad (2.D.1)$$

Then consider

$$\begin{aligned}
 \psi(F) &= \nabla\psi(F)F \\
 &= \sum_{\omega \in \Omega} \psi'(F, \mathbb{1}_{\omega})F(\omega) \\
 &= \sum_{i=1}^n \psi'(F, \mathbb{1}_{\sigma(i)})F(\sigma(i)), \quad (2.D.2)
 \end{aligned}$$

where the third inequality is just a reordering of the sum according to the payoff size of the states in F . Note that we have $\mathbb{1}_{\sigma(i)} = \mathbb{1}_{O(i)} - \mathbb{1}_{O(i-1)}$ and thus

$$\begin{aligned}
 \psi'(F, \mathbb{1}_{\sigma(i)}) &= \psi'(F, \mathbb{1}_{O(i)} - \mathbb{1}_{O(i-1)}) \\
 &= \psi'(F, \mathbb{1}_{O(i)}) - \psi'(F, \mathbb{1}_{O(i-1)})
 \end{aligned}$$

by homogeneity of ψ' in the second component. Let us now compute the Gâteaux derivative at F in the direction of $\mathbb{1}_{O(m)}$ for some $m \in \Omega$ by using the fact that F and $\mathbb{1}_{O(m)}$ are comonotonic:

$$\begin{aligned}\psi'(F, \mathbb{1}_{O(m)}) &= \lim_{\alpha \downarrow 0} \frac{\psi(F + \alpha \mathbb{1}_{O(m)}) - \psi(F)}{\alpha} \\ &= \lim_{\alpha \downarrow 0} \frac{\psi(F) + \alpha \psi(\mathbb{1}_{O(m)}) - \psi(F)}{\alpha} \\ &= \psi(\mathbb{1}_{O(m)}),\end{aligned}$$

where the second equality follows from Proposition 4.7 (Comonotonic Additivity) and the fact that F and $\alpha \mathbb{1}_{O(m)}$ are comonotonic and thus

$$\psi'(F, \mathbb{1}_{\sigma(i)}) = \psi(\mathbb{1}_{O(i)}) - \psi(\mathbb{1}_{O(i-1)}). \quad (2.D.3)$$

By combining (2.D.2), (2.D.3), and the definition of the capacity v we obtain (2.D.1). Hence, $U(f) = \psi(F) = \int_{\Omega} F dv$.

Next we show that v is indeed a capacity. We have $v(\emptyset) = 0$ by definition and $v(\Omega) = 1$ by the fact that $v(\Omega) = v(O(n)) = \psi(1_{\Omega}) = 1$. Moreover, we have to show that for any $X, Y \subseteq \Omega$ such that $Y \subset X$ we have $v(Y) \leq v(X)$. Since $v(Z) = \psi(\mathbb{1}_Z)$ for every $Z \subseteq \Omega$, it follows by monotonicity of ψ that $v(Y) \leq v(X)$.

Finally, we have to show that v is unique. Suppose there is another capacity w such that $U(f) = \int_{\Omega} F dw$ for all $f \in \mathcal{F}$ with $F = u \circ f$. Then take $F := \mathbb{1}_X$ for some $X \subseteq \Omega$ and let $k := |X|$. Note that $X = \{\sigma(i), \dots, \sigma(k)\}$. Then compute $U(f)$:

$$\begin{aligned}U(f) &= \int_{\Omega} F dv \\ &= \sum_{i=1}^n F(\sigma(i))(v(O(i)) - v(O(i-1))) \\ &= (v(\{\sigma(1)\}) - v(\emptyset)) + (v(\{\sigma(1), \sigma(2)\}) - v(\{\sigma(1)\})) + \dots + \\ &\quad (v(\{\sigma(1), \dots, \sigma(k)\}) - v(\{\sigma(1), \dots, \sigma(k-1)\})) + 0 \\ &= v(\{\sigma(1), \dots, \sigma(k)\}) - v(\emptyset) \\ &= v(X).\end{aligned}$$

By the same computation we have $U(f) = w(X)$ for every $X \subseteq \Omega$. Thus it must

hold that $v(X) = w(X)$ for all $X \subseteq \Omega$. \square

Proof of Proposition 4.9. Let $h := \alpha g + (1 - \alpha)\langle p \rangle$ where $p \in \Delta(X)$ is such that $u(p) = 0$. Take $q \in \Delta(X)$ such that $\langle q \rangle \sim h$. Since we have Full Independence, $\frac{1}{2}f + \frac{1}{2}h \sim \frac{1}{2}f + \frac{1}{2}\langle q \rangle$. Again we use Lemma 4.3 to show that $\psi(F + \alpha G) = \psi(F) + \alpha\psi(G)$ for any $\alpha \in (0, 1)$. By homogeneity of ψ we have that $\psi(F + \alpha G) = \psi(F) + \alpha\psi(G)$ holds for any $\alpha > 0$ (same argument as in the proof of Proposition 4.7). \square

2.E Appendix - Proofs Section 3.4

Proof of Proposition 4.12. Fix a $\alpha \in (0, 1]$ and choose $f \in \mathcal{F}$ as well as $p, q \in \Delta(X)$ such that $p \succeq^* f(\omega) \succeq^* q$ for all $\omega \in \Omega$. Then for $s_0 \in \Delta(X)$ such that $u(s_0) = 0$ it holds by Axiom 5 that

$$\alpha p + (1 - \alpha)s_0 \succeq^* \alpha f(\omega) + (1 - \alpha)s_0 \succeq^* \alpha q + (1 - \alpha)s_0,$$

for all $\omega \in \Omega$. Moreover, by Axiom 3 there exists a $\lambda \in [0, 1]$ such that

$$\lambda(\alpha\langle p \rangle + (1 - \alpha)\langle s_0 \rangle) + (1 - \lambda)(\alpha\langle q \rangle + (1 - \alpha)\langle s_0 \rangle) \sim \alpha f + (1 - \alpha)\langle s_0 \rangle.$$

Let $r := \lambda\langle p \rangle + (1 - \lambda)\langle q \rangle$, then the above can be written as $\alpha r + (1 - \alpha)\langle s_0 \rangle \sim \alpha f + (1 - \alpha)\langle s_0 \rangle$. For every $k \in [-1, 1]$, let $s_k \in \Delta(X)$ be such that $u(s_k) = k$ and apply Axiom 11 to obtain $\alpha r + (1 - \alpha)\langle s_k \rangle \sim \alpha f + (1 - \alpha)\langle s_k \rangle$ and thus $\psi(u \circ (\alpha\langle r \rangle + (1 - \alpha)\langle s_k \rangle)) = \psi(u \circ (\alpha f + (1 - \alpha)\langle s_k \rangle))$. Observe that

$$\begin{aligned} \psi(u \circ (\alpha f + (1 - \alpha)\langle s_k \rangle)) &= \psi(\alpha u \circ f + (1 - \alpha)u \circ \langle s_k \rangle) \\ &= \psi(\alpha F + (1 - \alpha)k\mathbb{1}_\Omega), \end{aligned}$$

where $F = u \circ F$. Furthermore, note that

$$\begin{aligned}
 \psi(u \circ (\alpha \langle r \rangle + (1 - \alpha) \langle s_k \rangle)) &= U(\alpha \langle r \rangle + (1 - \alpha) \langle s_k \rangle) \\
 &= u(\alpha r + (1 - \alpha) s_k) \\
 &= \alpha u(r) + (1 - \alpha) u(s_k) \\
 &= \alpha u(r) + (1 - \alpha) u(s_0) + (1 - \alpha) u(s_k) \\
 &= u(\alpha r + (1 - \alpha) s_0) + (1 - \alpha) k \\
 &= U(\alpha \langle r \rangle + (1 - \alpha) \langle s_0 \rangle) + (1 - \alpha) k \\
 &= U(\alpha f + (1 - \alpha) \langle s_0 \rangle) + (1 - \alpha) k \\
 &= \psi(\alpha F + (1 - \alpha) 0 \mathbb{1}_\Omega) + (1 - \alpha) k \\
 &= \psi(\alpha F) + (1 - \alpha) k.
 \end{aligned}$$

Thus we have

$$\begin{aligned}
 \psi(\alpha F) + (1 - \alpha) k &= \psi(u \circ (\alpha r + (1 - \alpha) \langle s_k \rangle)) \\
 &= \psi(u \circ (\alpha f + (1 - \alpha) \langle s_k \rangle)) \\
 &= \psi(\alpha F + (1 - \alpha) k \mathbb{1}_\Omega).
 \end{aligned}$$

Now suppose we have $\psi(F + \alpha \mathbb{1}_\Omega)$ for some $F \in [-1, 1]^\Omega$ and $\alpha \in (0, 1]$ such that $F + \alpha \mathbb{1}_\Omega \in [-1, 1]^\Omega$. Then we have

$$\begin{aligned}
 \psi(F + \alpha \mathbb{1}_\Omega) &= \psi\left((1 - \alpha) \left(\frac{1}{(1 - \alpha)} F\right) + \alpha \mathbb{1}_\Omega\right) \\
 &= \psi\left((1 - \alpha) \left(\frac{1}{(1 - \alpha)} F\right)\right) + \alpha = \psi(F) + \alpha.
 \end{aligned}$$

Thus we have shown that $\psi(F + \alpha \mathbb{1}_\Omega) = \psi(F) + \alpha$ for all $\alpha \in (0, 1]$. Consider $\alpha > 1$ and let $F \in [-1, 1]^\Omega$ then we have to show that $\psi(F + \alpha \mathbb{1}_\Omega) = \psi(F) + \alpha$ for $F + \alpha \mathbb{1}_\Omega \in [-1, 1]^\Omega$. Note that $\alpha \leq 2$ since $\psi \in [-1, 1]^\Omega$. Define $G := F + \mathbb{1}_\Omega$ and thus

$$\begin{aligned}
 \psi(F + \alpha \mathbb{1}_\Omega) &= \psi(F + \mathbb{1}_\Omega + \alpha \mathbb{1}_\Omega - \mathbb{1}_\Omega) \\
 &= \psi(G + (\alpha - 1) \mathbb{1}_\Omega),
 \end{aligned}$$

where $(\alpha - 1) < 1$. Therefore, it holds that $\psi(G + (\alpha - 1) \mathbb{1}_\Omega) = \psi(G) + (\alpha - 1) = \psi(F + \mathbb{1}_\Omega) + (\alpha - 1) = \psi(F) + 1 + (\alpha - 1) = \psi(F) + \alpha$ for all $\alpha > 1$. \square

Proof of Theorem 4.7. First we show that $U(f) = \nabla\psi(\hat{F})F$, where $\hat{F} = F - sF$ for some $s \in (0,1)$. Take some act $f \in \mathcal{F}$ and $F \in [-1,1]^\Omega$ such that $u \circ f = F$ and note that $\nabla\psi(\hat{F})F = \psi'(F - sF, F)$ by Lemma 4.1. Now take $F, (-F) \in [-1,1]^\Omega$ and since $(F - F) \in [-1,1]^\Omega$, by Theorem 4.6 there exists $s \in (0,1)$ such that $\psi(F - F) - \psi(F) = \psi'(F - sF, -F)$. By homogeneity of ψ' in the second component it follows that $\psi'(F - sF, -F) = -\psi'(F - sF, F)$. Since $\psi(F - F) = \psi(0) = 0$ we have that $\psi(F) = \psi'(F - sF, F)$ and therefore $U(f) = \psi'(F - sF, F)$ and consequently $U(f) = \nabla\psi(\hat{F})F$ as shown above. \square

Bibliography

- Qamrul Hasan Ansari, CS Lalitha, and Monika Mehta. *Generalized Convexity, Nonsmooth Variational Inequalities, and Nonsmooth Optimization*. CRC Press, 2013.
- Francis J Anscombe and Robert J Aumann. A definition of subjective probability. *Annals of mathematical statistics*, pages 199–205, 1963.
- Walter Armbruster and Werner Böge. Bayesian game theory. *Game theory and related topics*, pages 17–28, 1979.
- Geir B Asheim. *The consistent preferences approach to deductive reasoning in games*, volume 37. Springer Science & Business Media, 2006.
- Robert Aumann and Adam Brandenburger. Epistemic conditions for nash equilibrium. *Econometrica*, pages 1161–1180, 1995.
- Robert J Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, pages 1–18, 1987.
- Christian W Bach and Andrés Perea. Utility proportional beliefs. *International Journal of Game Theory*, 43(4):881–902, 2014.
- Truman Bewley. Knightian decision theory. *Part I: Decisions in Economics and Finance*, 25, 2002.
- Mariana Blanco, Dirk Engelmann, Alexander K Koch, and Hans-Theo Normann. Belief elicitation in experiments: is there a hedging problem? *Experimental Economics*, 13(4):412–438, 2010.
- W. Böge and Th. Eisele. On solutions of bayesian games. *International Journal of Game Theory*, 8(4):193–215, 1979.

- Adam Brandenburger and Eddie Dekel. Rationalizability and correlated equilibria. *Econometrica: Journal of the Econometric Society*, pages 1391–1402, 1987.
- Adam Brandenburger and Eddie Dekel. The role of common knowledge assumptions in game theory. In F. Hahn, editor, *The economics of missing markets, information and games*, pages 46–61. Oxford University Press, 1989.
- Simone Cerreia-Vioglio, Paolo Ghirardato, Fabio Maccheroni, Massimo Marinacci, and Marciano Siniscalchi. Rational preferences under ambiguity. *Economic Theory*, 48(2-3):341–375, 2011.
- Yves Chabrilac and J-P Crouzeix. Continuity and differentiability properties of monotone real functions of several real variables. In *Nonlinear analysis and optimization*, pages 1–16. Springer, 1987.
- Juin-Kuan Chong, Colin F Camerer, and Teck-Hua Ho. Cognitive hierarchy: A limited thinking theory in games. *Experimental Business Research*, pages 203–228, 2005.
- Miguel A Costa-Gomes and Vincent P Crawford. Cognition and behavior in two-person guessing games: An experimental study. *The American economic review*, pages 1737–1768, 2006.
- Miguel A Costa-Gomes and Georg Weizsäcker. Stated beliefs and play in normal-form games. *The Review of Economic Studies*, 75(3):729–762, 2008.
- Miguel A Costa-Gomes, Vincent P Crawford, and Bruno Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, pages 1193–1235, 2001.
- Vincent P Crawford and Nagore Iriberry. Fatal attraction: Salience, naivete, and sophistication in experimental "hide-and-see" games. *The American Economic Review*, pages 1731–1750, 2007.
- Anna Dreber, Tore Ellingsen, Magnus Johannesson, and David G Rand. Do people care about social context? framing effects in dictator games. *Experimental Economics*, 16(3):349–371, 2013.
- Daniel Ellsberg. Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, pages 643–669, 1961.

- Urs Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178, 2007.
- Daniel E Fragiadakis, Daniel T Knoepfle, and Muriel Niederle. Identifying predictable players: Relating behavioral types and subjects with deterministic rules. *Unpublished manuscript*, 2013.
- Sotiris Georganas, Paul J Healy, and Roberto A Weber. On the persistence of strategic sophistication. 2014.
- Paolo Ghirardato, Fabio Maccheroni, and Massimo Marinacci. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(2):133–173, 2004.
- Itzhak Gilboa. *Theory of decision under uncertainty*, volume 1. Cambridge university press Cambridge, 2009.
- Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, 18(2):141–153, 1989.
- Jacob K Goeree and Charles A Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, pages 1402–1422, 2001.
- Jacob K Goeree, Charles A Holt, and Thomas R Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1): 97–113, 2003.
- Jacob K Goeree, Charles A Holt, and Thomas R Palfrey. Regular quantal response equilibrium. *Experimental Economics*, 8(4):347–367, 2005.
- Ben Greiner. An online recruitment system for economic experiments. 2004.
- John C Harsanyi. Games with incomplete information played by "bayesian players" part i, ii , iii. *Management Science*, 14(5):159–182, 320–334, 486–502, 1967-68.
- Daniel Kahneman. *Thinking, fast and slow*. Macmillan, 2011.
- Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, pages 1447–1498, 2006.
- Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6 – 38, 1995.

- Jean-François Mertens and Shmuel Zamir. Formulation of bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14(1): 1–29, 1985.
- Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326, 1995.
- John Nash. Equilibrium points in n-person games. *Proc. Nat. Acad. Sci. USA*, 36(1):48–49, 1950.
- John Nash. Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295, 1951.
- Christian T Nauerz. Beyond coincidence: the reasoning process underlying utility proportional beliefs process. *Working Paper*, 2015.
- Christian T Nauerz, Marion Collewet, and Frauke Meyer. Identifying strategic thinking in one-shot games. *Working Paper*, 2016.
- Theo Offerman and Asa B Palley. Losses in translation: an off-the-shelf method to recover probabilistic beliefs from loss-averse agents. *Experimental Economics*, pages 1–30, 2015.
- Theo Offerman, Joep Sonnemans, Gijs Van de Kuilen, and Peter P Wakker. A truth serum for non-bayesians: Correcting proper scoring rules for risk attitudes. *The Review of Economic Studies*, 76(4):1461–1489, 2009.
- Efe A Ok. *Real analysis with economic applications*, volume 10. Princeton University Press, 2007.
- Asa B Palley. Great expectations: Prospect theory with a consistent reference point. *Available at SSRN 2313851*, 2013.
- Andrés Perea. A one-person doxastic characterization of nash strategies. *Synthese*, 158(2):251–271, 2007.
- Andrés Perea. *Epistemic game theory*. Cambridge University Press, 2012.
- Ben Polak. Epistemic conditions for nash equilibrium, and common knowledge of rationality. *Econometrica*, 67(3):673–676, 1999.
- Pedro Rey-Biel. Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior*, 65(2):572–585, 2009.

- Ariel Rubinstein. "economics and psychology"? the case of hyperbolic discounting*. *International Economic Review*, 44(4):1207–1216, 2003.
- Leonard J Savage. *The foundations of statistics*. Courier Corporation, 1972.
- David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, pages 571–587, 1989.
- R Selten and S Krischker. Comparison of two theories for characteristic function experiments. In *Aspirations Levels in Bargaining and Economics Decision Making*, pages 259–264. Springer, 1983.
- Reinhard Selten. Properties of a measure of predictive success. *Mathematical Social Sciences*, 21(2):153–167, 1991.
- Dale O Stahl and Paul W Wilson. Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3):309–327, 1994.
- Dale O Stahl and Paul W Wilson. On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.
- Keith E Stanovich, Richard F West, et al. Individual differences in reasoning: Implications for the rationality debate? *Behavioral and brain sciences*, 23(5):645–665, 2000.
- Tommy C Tan and Sérgio R Werlang. The bayesian foundations of solution concepts of games. *Journal of Economic Theory*, 45(2):370–391, 1988.
- John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1944.
- Georg Weizsäcker. Ignoring the rationality of others: evidence from experimental normal-form games. *Games and Economic Behavior*, 44(1):145–171, 2003.

Valorization

Knowledge valorization is the process of translating academic knowledge into social benefits. Scientific output and impact are difficult to measure and even harder to predict. This is particularly the case in the short-term. The history of science is full of remarkable ideas, which only gained popularity decades after they were conceived. The problem becomes even worse when we consider fundamental research. Clearly, fundamental insights provide the foundation for many applications but often applications only arise long after a fundamental problem has been solved. So at the time the research was conducted it might have been considered useless, and years later as indispensable. To give the reader an idea of how the insights of this thesis can benefit society, this addendum will outline some conceivable applications.

The research presented in this thesis is concerned with fundamental problems in game and decision theory. No economic or industrial applications are discussed making immediate valorization hard to detect. Therefore, we will discuss the game theoretic findings and the decision-theoretic separately to highlight its respective contributions better.

Classically, game theoretic problems are analyzed using the famous Nash equilibrium. Nash equilibrium aims at finding combinations of choices under which none of the involved persons has an incentive to deviate from her choice. The concept can be very useful when applied in the right context but used in the wrong context it can be of little help. Situations in which a strategic interaction is repeated several times might be described accurately by Nash equilibrium. In this situation, people have several chances to learn how their opponents acted in the past and can adapt their actions accordingly. Thus they might end up choosing such that in the following rounds there is no incentive to choose differently as

the opponent choose a certain action repeatedly. If an analyst is interested in such a setting, Nash equilibrium can be an excellent tool. Many economic decisions, however, do not occur repeatedly with the same set of agents involved. In these cases, Nash equilibrium does usually not perform well in explaining peoples' choices. Nevertheless, Nash equilibrium is often used to analyzed non-repeated settings. The most important large scale applications of Nash equilibrium lie in the field of antitrust cases. Antitrust considerations usually involve billions of dollars and possibly drastic consequences for society in the long-term. Of course, there is more to an antitrust decision than a game theoretic model. However, often game theoretic models are used as an underlying intuition of what might happen, for example, in the case of a merger. This might not affect clear cut cases but possibly border cases.

Experimental evidence in this thesis shows that Nash equilibrium or even its weaker version, quantal response equilibrium, do not perform well in describing and therefore by implication in predicting peoples' choices. Consequently, this research shows the need to adapt the models used in practice. Practitioners should have a close look at the outcome of this experiment and related results in the literature to adapt their models if needed. Fortunately, the models that perform well are also easy to implement.

In antitrust cases, there might be several actors to be considered. First of all, the companies in questions and their respective competitors. Secondly, we have to consider the customers. In a business-to-consumer market, it is possible that companies are more sophisticated in their decision making than consumers. However, the consumers' behavior has a significant impact on the outcome of antitrust decisions. The first project presented in this thesis shows how game theoretic models can be tested more extensively. Therefore, those insights might help to build more realistic game theoretic models allowing more accurate predictions of resulting behavior and by implication better antitrust decisions.

Next to these high impact applications there are also countless applications in everyday life and business. One of the most impactful scenarios for society is so-called public good games. These games are characterized by the fact that the outcome of the group is maximized if everybody participates in producing the public good, e.g., not overfishing a certain species of fish in a given region. The challenge, however, comes from the fact that it is always better for any participant to deviate. In the fishing example, this means that if all participants stuck to their respective contingent, it would be better for the individual fisher to fish

more (given that there is no punishment). These kinds of interactions occur in many fashions, e.g., the emission of CO₂, use of grounds for farming, maintenance of public infrastructure, and so on. Hence, it is vital for policy makers to have a clear understanding of the participants' behavior to different incentive structures. Even though, Nash equilibrium might be better suited to describe long-term outcomes, a better understanding of initial reactions is vital. Nash equilibria are often not unique, and the analyst can therefore not predict in which pair of choices will prevail. Experimental evidence shows that initial choices can have an impact on the path towards equilibrium. Therefore, the insights of this thesis can help to understand how people will initially react to a new policy and what equilibrium they are likely to end up in. Consequently, better-informed decisions can have large scale impact on society many important settings.

Game theory is also likely to provide intuition for many small scale business decisions. Almost all business degrees include at least one class dedicated to game theory. These classes will not make the students experts in game theoretic analysis. However, they will certainly shape their approach to strategic interactions. Again, many business decisions are not repeated and can therefore not be accurately described using Nash equilibrium. Taking the insights of this thesis and those of the related literature to implement the more accurate models in the curriculum can also have an indirect effect on society through many small scale decisions.

To sum up, the impact of game theoretic findings on society, recall that situations requiring game theoretic thinking are abundant and appear in large scale applications as well as in small-scale applications. The fundamental research presented in this thesis might not directly contribute to a concrete problem, but it will certainly help future generation to make better strategic decisions in a variety of contexts that have a strong impact on society.

The impact of the decision-theoretic research in this thesis is also not immediately linked to societal problems. Decision-theoretic models build the foundation of almost all economic models. They describe how an economic agent chooses given a set of possible choices. This feature appears in many macroeconomic models as well as in microeconomic models. Central banks use macroeconomic models to predict countries GDP, unemployment rate, and many more variables that are taken into account when considering the supply of money needed to support the economy. There are various macroeconomic models. Some of the more sophisticated ones also take into account the uncertainty economic

agents face. This is where the research in this thesis becomes relevant. The uncertainty economic agents encounter can often not be described with an objective probability. In these cases, the standard subjective expected utility model, which underlies most macroeconomic models, has shown to exclude plausible choices. Thus, macroeconomic models may fail to represent important contingencies and might arrive eventually at incorrect conclusions. Again, it is unlikely that applying an alternative to subjective expected utility will change stable outcomes, but there might be borderline cases where it can have an impact on the predictions. Macroeconomics has yet to fully adopt alternative models of decision making under uncertainty. One of the obstacles is that there is a variety of models and no objective favorite. The research in this thesis shows, however, how most models of decision making under uncertainty are linked—a matter that was previously not well understood. This will help practitioners to make more reliable predictions.

As mentioned earlier, decision theory is the basis for most economic models. Therefore, it also plays a major role in many microeconomic settings with uncertainty. Microeconomics can be used to describe a huge variety of social interactions. For example, it can be used to describe how a person should allocate her time. An example of a high impact decision is the education a person should obtain. Does it pay off in terms of person's utility to get a bachelor degree, or is it better to get a master degree, or maybe a more elementary degree will be beneficiary. This kind of decision depends on the demands of the labor market as well as the person's preferences for money and leisure time. Uncertainty arises in this context through the difficulty of determining objective probabilities for the outcomes of the considered paths induced by the obtained education. Another example is the optimal usage of utilities like water, gas, and electricity. Often the prices increase progressively, which makes choosing optimally more challenging than with linear pricing schemes. Again, a person has to predict its future use of resources to choose optimally, and the prediction is a non-objective probability distribution about future events. Microeconomics covers such a broad range of phenomena that it is impossible to give an exhaustive list of possible implications of the research conducted in this thesis. It should be clear, however, that even though no immediate applications are discussed, the potential impact on society can be significant.

In conclusion, this thesis is concerned with fundamental research in decision and game theory. Basic building blocks of economic models are assessed and

brought closer to peoples' actual behavior. These insights can inform practitioners in policy making and business to make more accurate predictions of peoples' behavior and therefore benefit society. Even though fundamental insights do not translate directly into recommendations, they provide insights which will improve the quality of many large and small scale decisions.

Nederlandse samenvatting

Dit proefschrift onderzoekt redeneren over het redeneringsproces dat hoort bij (strategische) besluitvorming onder onzekerheid. Onzekerheid beschrijft situaties waarin de beslisser geen objectieve kans van optreden kan toekennen aan een bepaalde gebeurtenis. Twee klassieke voorbeelden die gebeurtenissen. Hierin onderscheidt het verschil tussen zich van situaties met risico en onzekerheid illustreren zijn respectievelijk . Denk bijvoorbeeld aan loten waar waarbij de kans op winnen/winstkans objectief bepaald kan worden (risico) en weddenschappen op bij paardenrennen waarwaarbij deze kans niet objectief te bepalen is. Ten opzichte van (onzekerheid). Zeker over besluitvorming onder risico (loten) valt er over besluitvorming onder onzekerheid valt nog veel leren.

Beslissingsproblemen met onzekerheid kunnen onder verschillende omstandigheden ontstaan. Speltheorie onderzoekt strategisch de strategische besluitvorming waar dewaardoor het nut van de ene speler afhankelijk is van de keuze van een ander. Een strategische besluitbeslissing in zo'n geval is tekenen men opvatten beschouwen als een besluitbeslissing onder onzekerheid. In plaats van een subjectieve waarschijnlijkheid toerekenen kans toe te kennen aan de gebeurtenis dat een bepaalde bepaald paard de wedstrijd wint, kijken we naar de waarschijnlijkheid dat een andere speler een bepaalde keuze maakt. Zodra de persoon 's Zodra bekend is welke kansen een persoon toekent aan de keuzes van een tegenstander, kunnen we de besluitvorming van deze persoon opvatten als beslissen waarschijnlijkheidstoekenning aan de tegenstander 's keuzes bekend is, staat een strategische besluit gelijk aan een besluit onder onzekerheid.

Gezien het standaard besluitvormingsmodel Het standaard beslissingsmodel onder onzekerheid goed presteert relatief goed in de meeste contexten dient de

focus tesituaties. Derhalve verschuiven we onze aandacht naar het begrijpen van het menselijke redeneringsproces in een besluit probleem. Een beter begrip van beslissingsprobleem. Als we dit redeneringsproces helpt in hetbeter begrijpen, dan kunnen we ook een beter inzicht verwerven van inzichten in de manier waarop deen persoon de subjectieve kansen toekent, en hoe deze persoon uiteindelijk haar beslissing neemt. waarschijnlijkheidstoekenning is gevormd en dus in het bepalen van de uiteindelijke besluit. Een bevatting van de rol van de waarschijnlijkheidstoekenningDit, op haar beurt, vereist een beter begrip van het basisprincipe van het standaardmodel van besluitvormingbeslissen onder onzekerheid: het subjectief verwachte nut. Subjectief verwachte nut veronderstelt dat hetEen individu een waarschijnlijkheid toekentmet subjectief verwacht nut kent aan ieder mogelijke gebeurtenis. een subjectieve kans toe. In het voorbeeld van weddenschappen opbij paardenrennen betekent dit dat de speler een winningskanswinstkans toekent aan iedere paard. Gegeven deze waarschijnlijkheidstoekenning wordt hetHet nut van elke uitkomst van een loterij gewogen naarwordt vermenigvuldigd met de kans dat die zich zal voordoen. Als de winningskanswinstkans van paard A 80% is, en het loterij lot voor dit paard 20 euro zou uitkeren als paard A wint, dan is de verwachte waarde van het lot 16 euro. Op deze manier bepaalt deZo bepaalt de beslisser zijn keuze op basis van de subjectieve kansen die hij toekent aan de verschillende gebeurtenissen. waarschijnlijkheidstoekenning van Binnen de beslisser zijn keuze.

In speltheorie, de studie van een vakgebied dat strategische interacties, zijn er modellen die tussen mensen bestudeert, kijken we naar aannemelijke waarschijnlijkheidstoekenningen beschrijven kansen die een persoon kan hebben. Van-toekennen aan de keuzes van zijn tegenstanders. Zulke kansen worden vanaf nu af aan noemen we een individu's "beliefs" genoemd. waarschijnlijkheidstoekenningen in een situatie zijn overtuiging. Bach en Perea (2014) ontwerpen en speltheoretischhebben een spel-theoretisch model ontwikkeld waarin de overtuigingen van kans die een speler toekent aan een keuze van zijn tegenstander evenredig dienendient te zijn aan het subjectief verwachte nut geïnduceerd door de overtuigingen van zijndat deze keuze oplevert voor de tegenstander. tegen-speler over zichzelf. Dit idee is het primitieve concept van het de basis voor hun model. Een ander belangrijk onderdeel vaak voorkomend in speltheorie De volgende stap in hun model is de overtuiging van een speler dat een ander zijn primitief deelt. Met andere woorden,spelers aannemen dat die anderhun tegenstanders ook overtuigingen heeft die evenredig zijn aan het subjectief verwachte nut

geïnduceerd door een overtuiging over de eerste persoon. Daarenboven wordt er veronderstelt op een soortgelijke manier redeneren. En dat een persoon geloofthun tegenstanders denken dat de ander overtuigingen heeft die evenredig zijn aan subjectief verwachte nutandere spelers ook op een soortgelijke manier redeneren, enzovoort. Onder de, veronderstelling dat deze primitieve conditie van toepassing is op alle orden van overtuiging, kunnen alle aannemelijke overtuigingen onder hetAl deze aannames samen leiden tot een epistemisch concept. Bach en Perea hebben een algoritme ontworpen dat uiteindelijk leidt tot de beliefs die mogelijk zijn in dit epistemische concept bepaald worden.. Dit zijn echter epistemische condities die beperkingen opleggen aan mogelijke overtuigingen van spelers op iedere orde van overtuiging. Om deze overtuigingen te bepalen ontwerpen Bach en Perea een algoritme waarmee overtuigingenverwijdert iteratief beliefs die inconsistent zijn met het redeneringsconceptepistemische concept. iteratief worden verwijderd. Het nadeel van dit algoritme is dat de gesuggereerde berekeningen onwaarschijnlijk op een soortgelijke wijzein iedere stap waarschijnlijk niet als zodanig door mensen uitgevoerd zullen worden. Dit In dit proefschrift bezorgt een alternatieve benadering om overtuigingen consistent met het concept onder beschouwing af te bakenen in de vorm vanlaat ik zien dat het algoritme ook omgevormd kan worden tot een iteratieve redeneringsprocedure. De voorgesteldeIn iedere stap van deze procedure past in iedere stap de overtuiging persoon haar beliefs aan. Het proces is tweedelig: (1) in de initiële overtuiging is bestaat uit twee delen: In het begin is de kans die een persoon toekent aan een keuze zo goed als van de tegenstander evenredig aan het gemiddelde nut suggereert, envan deze kans. Later past de persoon deze kansen aan, door zichzelf in de schoenen van te tegenstander te plaatsen. (2) een iteratieve redeneringsprocedure waar de drijfveren van de speler en zijn tegenstander iteratief in beschouwing worden genomen om de vorige overtuiging aan te passen. Deze visie op het conceptinterpretatie is niet alleen nuttig omdat het aannemelijker is dat mensen op deze wijze volgens het concept zouden redeneren, maar ook omdat het testen van een redelijke beschrijving kan zijn van hoe mensen in spel-theoretische situaties redeneren. Bovendien is het met deze interpretatie mogelijk om kenmerken van hetde procedure mogelijk maakt.te testen. Dit laatste is onmogelijk in het geval van hetniet mogelijk met de iteratieve procedure van Bach en Perea. Een toets van deze kenmerken laatGedragsexperimenten laten zien dat er inderdaad mensen zijn die volgens dit concept lijken te redeneren. DeDeze resultaten verwerven een geven ons derhalve inzicht in de manier

waarop sommige mensen redeneren onder strategische onzekerheid proberen op te lossen. Met gebruik van deze inzichten kankunnen we het gedrag van een subpopulatiegroep mensen (die op deze manier redeneren) worden voorspelt-voorspellen, en kunnen beleidsaanbevelingen in eenwe aanbevelingen doen aan beleidsmakers die met deze strategische situatie vervolgenssituaties geconfronteerd worden. gemaakt worden.

Om beleidsaanbevelingengoede aanbevelingen te kunnen makendoen is een begrip van het gedrag van een groter deel van de populatie echter noodzakelijk. Derhalve presenteert om het gedrag van een grotere groep mensen te kunnen voorspellen. In dit proefschrift een experiment om diverspresenteer ik een gedragsexperiment dat verschillende aannemelijke modellen van strategisch redeneren te testen. Als mensen werkelijk op diverse manieren redeneren in strategische situaties zal het test. Het experiment de relevantetoont aan welke redeneringsmodellen onthullen en voor iederrelevant zijn, en identificeert voor ieder van deze modellen de groep mensen die volgens dit model redeneert. dat deel van de populatie die volgens het model redeneert identificeren. Dit is niet het eerste experiment met dit doel. Dit is nietechter wel het eerste experiment dat de verschillende wijze waarop mensen strategisch denken benaderen probeert te identificeren. Niettemin is dit wel de eerste keer dat er gebruik wordt gemaakt van gedragsexperiment in deze categorie waarin de strategische beslissingsproblemen met weinigniet tot unieke keuzes leiden met conventionele iteratieve procedures. redeneringsstructuur (waar geen één keuze altijd optimaal is los van de keuze van de tegenstander). Verder iszijn de methode vernieuwend gezien er gebruik gemaakt wordt van zwakkere aannames in dit experiment zwakker dan men voorheen nodig geacht. achtte.

Mitsdien bevorderen de resultaten van het experiment de vertrouwen in het onderzoeksveld gezien de modellen goed te lijken presteren onder veeleisende omstandigheden (in strategische beslissingsproblemen met weinig redeneringsstructuur). Resultaten suggereren dat het k-level redeneringsmodel en het concept van Bach en Perea de redenering van ongeveer 80% van de deelnemers beschrijft. Het level-k redeneringsmodel veronderstelt dat de belief van een level-0 persoon gelijke kansen toekent aan alle keuzes van zijn tegenstander. Een level-1 persoon maakt een optimale keuze gegeven zijn belief dat zijn tegenstander een level-0 persoon is. Een level-2 persoon maakt een optimale keuze gegeven zijn belief dat zijn tegenstander een level-1 persoon is, enzovoort. Dit concept ligt erg dicht bij hoe men het menselijke redeneringsproces voor kan stellen en is daarom

zeer succesvol geweest. Ook zijn er kenmerken van het level-k model die getest kunnen worden. Spelers die volgens het level-k model redeneren voldoen ook werkelijk aan deze kenmerken.

De twee bovengenoemde delen van dit proefschrift gaan over het oplossen van onzekerheid in situaties van strategische interactie. Echter, onzekerheid kan optreden in simpelere situaties. In een urn met 50 rode en 50 blauwe ballen is de kans op trekking van een rode bal is 50%. Denk aan een urn met 100 ballen waar iedere bal rood of blauw kan zijn. In dit geval is er geen vanzelfsprekende manier om een objectieve kans toe te kennen omdat de verdeling van kleuren onbekend is. In zo'n geval kan het klassieke subjectief verwachte nut model plausibele keuzes uitsluiten (zie bijvoorbeeld het Ellsberg paradox). Diverse modellen in de literatuur probeerde dit probleem op te lossen, vaak door het toestaan van meerdere kans verdelingen per speler. Het verwachte nut van een keuze is dan het minimaal verwachte nut gegeven alle kans verdelingen van de speler. Gezien de verhouding tussen deze modellen onduidelijk is, presenteert dit proefschrift een manier om ze te verenigen en daarmee nieuwe inzichten verwerven relatie tussen de verschillende benaderingen. Deze inzichten zorgen, op hun beurt, voor een beter begrip van de wijze waarop mensen keuzes maken onder dit soort onzekerheid.

Ten slotte onderzoekt dit proefschrift het gedrag van mensen in beslissingsproblemen onder onzekerheid. Onzekerheid beschrijft situaties waarin de beslisser geen objectieve kans van optredencan toekennen aan eende gebeurtenis toekennen. Twee klassieke voorbeelden diegebeurtenissen onder beschouwing. Het standard model van subjectief verwachte nut presteert relatief goed in de veel contexten dient de focus tesituaties. In dit soort situaties is de manier waarop spelers hun kans verdeling vormen een zeer belangrijke vraag. Van bijzonder belang is het geval van strategische beslissingen. Zoals dit proefschrift laat zien is iteratief denken over mens eigen nut en dat van de tegenstander een veelgebruikte benadering. En hoewel, iteratief redenering een belangrijke rol speelt, lijken er ten minste twee verschillende redeneringsmodellen van belang te zijn. Om rekening te houden met alle redelijke keuzes onder onzekerheid dient het standaard verwachte nut model echter te veranderen. Dit is te bereiken op verschillende manieren die verenigd zijn in dit proefschrift.

Curriculum Vitae

Christian Tobias Nauerz was born on September 6, 1987, in Augsburg, Germany. In 2007, he received his Abitur from the Lichtenberg Schule Darmstadt. After finishing his civil service, he studied business administration at the University of Eichstätt-Ingolstadt and obtained the bachelor degree in 2011. In 2013, he graduated from the master program in economic and financial research with distinction.

The same year Christian started the Ph.D. program with Andrés Perea and Elias Tsakas as daily supervisors. From January 2016 to July 2016, he spent six months visiting the University of California, Berkeley. His research is presented in this thesis. During the Ph.D. program, Christian presented his research at several international conferences.