

Incorporating Ignorance within Game Theory: An Imprecise Probability Approach

by

Bernard Fares

Submitted to the School of Computer Science and Statistics
in partial fulfilment of the requirements for the degree of

Doctor of Philosophy in Statistics

at the

TRINITY COLLEGE DUBLIN

2023

© Trinity College Dublin 2023. All rights reserved.

School of Computer Science and Statistics

Thesis Supervisor:

Mimi Zhang, Associate Professor

External Reviewer:

Frank P.A. Coolen, Professor

Dedicated to my mother and the memory of my father.

Incorporating Ignorance within Game Theory: An Imprecise Probability Approach

by

Bernard Fares

Submitted to the School of Computer Science and Statistics
on November 18, 2022, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Statistics

Abstract

Ignorance within non-cooperative games, reflected as a player's uncertain preferences towards a game's outcome, is examined from a probabilistic point of view. This topic has had scarce treatment in the literature, which emphasises exogenous uncertainties caused by other players or nature and not by players themselves. That is primarily because a player's endogenous uncertainty over an outcome poses significant challenges and complex sequences of reciprocal expectations. Therefore, it is often ignored, and preferences are either assumed from a continuous domain or set using introspection.

Decisions under ignorance could be optimised by permitting a player to compute rational strategies with respect to elicited lower and upper expectations of an uncertain outcome, allowing them to update these strategies when new observations are available, and helping them assess the impact and value of acquired information. Therefore, this dissertation aims to develop a complete framework for decision optimisation within strategic settings that include uncertainty. We explore a solution concept based on recent research in imprecise probabilities and de Finetti's approach to defining subjective probabilities, which utilises bets to assess beliefs.

An in-depth literature review of game theory and imprecise probabilities is provided, focusing on existing normative theories and their plausible generalisations. The motivation behind a solution permitting ignorance is presented, and foundational issues related to existing approaches are argued. Afterwards, we introduce a framework that allows a risk-neutral player with constant marginal utilities for money to incorporate and dynamically learn about uncertain outcomes. This

framework is then generalised to cover risk-averse players whose marginal utilities across outcomes are state-dependent.

The resulting framework is proposed as a possible solution to the problem of utility induction and decision-making in game-theoretic settings that include uncertainty. It is analysed and demonstrated through motivating examples modified to include uncertainty. Each example's correlated equilibria's convex polytope is computed and compared to its uncertainty-free equivalent. Exceptional cases such as extreme ignorance are also examined and assessed through a Monte Carlo simulation where we demonstrate that, in repeated games, vacuous lower and upper previsions converge to one linear value that reflects the *true* expected preference over the uncertain outcome.

Moreover, inadequate value of information under uncertainty is considered, and a model to assess the impact of information patterns on strategic interactions is proposed. This model enables a player to compute their expected and actual values of a piece of information with respect to a Pareto-efficient strategy. We showcase it within a game that includes uncertainty by applying utility diagnostics to two types of players, pessimistic and optimistic.

Finally, since the foundations of the normative game theory introduced by Von Neumann and Morgenstern assume that all outcomes are known, the consistency of its axiomatic rules under ignorance is reviewed. We show that uncertainty can alter relevant games' zero-sum and symmetry properties and propose an approach to force these properties.

Thesis Supervisor: Mimi Zhang
Title: Associate Professor

Acknowledgments

First and foremost, I would like to express my great appreciation to Dr Brett Houlding. Although not with us anymore, he continues to inspire me with his work and contribution to the decision theory realm. As my previous supervisor, Brett's guidance and expert advice were invaluable and significantly influenced my work. It was an honour to work with a person who achieved so much in his life and yet remained modest enough to sacrifice his time to help students.

A special thank you to my thesis supervisor, Dr Mimi Zhang, for her continuous help and supervision. Her guidance and feedback helped me develop my skills and improve the quality of my deliverables.

I would like to express my gratitude to Citibank for partially funding my PhD, backing me in my journey, and allowing me to pursue my goal of becoming a researcher in statistics.

To conclude, a big thanks to my family, friends, and partner Leanne for all the unconditional love and support!

Contents

List of Figures	10
List of Tables	12
1 Introduction	15
1.1 Utility Theory	15
1.2 Problem Description	16
1.3 Outline of Thesis	18
2 Review of Game Theory	19
2.1 Game Theory	19
2.2 Information in a Game	23
2.3 Strategic Form of Games	24
2.4 Risk Attitudes	26
2.5 Rational Behaviour	29
2.6 The Strategy	30
3 Imprecise Probabilities and Related Literature	39
3.1 Imprecise Probabilities	39
3.2 Revealed-Rules Matrix	43
3.3 Non-parametric Utility Updating	50
4 Suggested Model Under Ignorance	53
4.1 Motivation	54

4.2	Enhanced Revealed-Rules Matrix	58
4.3	Rationality Requirements on Lower and Upper Previsions	62
4.4	First Assessment and Refinement	63
4.5	Dynamic Updating	69
4.6	Risk Aversion	74
4.7	Game Expectation Under Uncertainty	79
4.8	Outcome's Expected Net Payoff	80
4.9	Example	82
5	Special Cases	97
5.1	Extreme Ignorance	97
5.2	Multiple Sources of Ignorance	101
5.3	Ignorance within Zero-Sum Games	103
5.4	Ignorance within Symmetric Games	107
6	Utility Diagnostics Under	
	Uncertainty	111
6.1	Value of Information	111
6.2	Example	117
7	Discussion and Conclusions	129
7.1	Discussion	129
7.2	Conclusions	131
	List of Symbols	134
	Bibliography	139

List of Figures

2-1	‘Battle of the sexes’ polytope and set of all joint independent probability distributions between players - Intersections TL and BR are pure Nash equilibria, and intersection $(2/3 \text{ T}, 1/3 \text{ B}) \times (1/3 \text{ L}, 2/3 \text{ R})$ is a mixed Nash equilibrium.	37
3-1	‘Battle of the sexes’ with risk-averse players - Polytope and set of all joint independent probability distributions between players. . . .	49
4-1	Example’s credal set.	85
4-2	Modified version of ‘battle of the sexes’ with risk-neutral players - First assessment polytope.	88
4-3	Modified version of ‘battle of the sexes’ with risk-averse players - First assessment polytope.	92
5-1	Modified version of ‘matching pennies’ - Average lower and upper previsions generated using 1000 simulations with 200 plays each . . .	99
5-2	Modified version of ‘matching pennies’ - Players’ average game expectations generated using 1000 simulations, with 200 plays each. . .	100

6-1	Player one's expected payoff and initial assessment of gamble f 's lower prevision, given the probability of a riot. The expected payoff is based on a Pareto efficient correlated strategy and on the assumption that player one is extremely pessimistic.	122
6-2	Player one's expected payoff and initial assessment of gamble f 's upper prevision, given the probability of a riot. The expected payoff is based on a Pareto efficient correlated strategy and on the assumption that player one is extremely pessimistic.	124
6-3	Expected value of sample information of both a pessimistic and optimistic player one, given the probability of a riot.	126
6-4	Expected payoff of a pessimistic and optimistic player one given the <i>ex-ante</i> Pareto efficient strategy and given that a riot starts.	127
6-5	Value of a 'riot' for a pessimistic and optimistic player.	128

List of Tables

2.1	‘Battle of the sexes’ - Payoff matrix.	28
2.2	‘Battle of the sexes’ with risk-neutral players - Polytope vertices. . .	38
3.1	‘Battle of the sexes’ - Revealed-rules matrix \mathbf{M}	45
3.2	‘Battle of the sexes’ with risk-averse players - Payoff matrix, de- nominated in utility units.	47
3.3	‘Battle of the sexes’ with risk-averse players - Revealed-rules matrix M^*	48
3.4	‘Battle of the sexes’ with risk-averse players - Polytope vertices. . .	48
4.1	Modified version of ‘battle of the sexes’ with risk-neutral players - Payoff matrix.	59
4.2	Modified version of ‘battle of the sexes’ with risk-neutral players - GRR matrix with uncertain payoffs modelled as gambles.	60
4.3	Modified version of ‘battle of the sexes’ - Enhanced GRR matrix, where gamble f is replaced by its lower and upper previsions. . . .	61
4.4	Modified version of ‘battle of the sexes’ with risk-averse players - Payoff matrix.	75

4.5	Modified version of ‘battle of the sexes’ with risk-averse players - Enhanced GRR matrix, where gamble f ’s utility is replaced by its lower and upper previsions.	76
4.6	Modified version of ‘battle of the sexes’ - Bought payoff vectors. . .	82
4.7	Modified version of ‘battle of the sexes’ - Outcomes net payoffs. . .	82
4.8	Modified version of ‘battle of the sexes’ with risk-neutral players - Resulting enhanced GRR matrix.	86
4.9	Modified version of ‘battle of the sexes’ with risk-neutral players - First assessment vertices.	87
4.10	Modified version of ‘battle of the sexes’ with risk-averse players - Resulting enhanced GRR matrix.	90
4.11	Modified version of ‘battle of the sexes’ with risk-averse players - First assessment vertices.	93
4.12	Dynamic updating applied to three different scenarios.	93
5.1	Modified version of ‘matching pennies’ - Payoff matrix.	98
5.2	Modified version of ‘matching pennies’ with risk-neutral players - Enhanced GRR matrix.	104
5.3	Modified version of ‘matching pennies’ with risk-neutral players - Finalised enhanced GRR matrix.	104
5.4	Modified version of ‘prisoner’s dilemma’ with risk-neutral players - Payoff matrix.	108

Introduction

This chapter describes the problem examined in this study. We start by giving an overview of *utility theory* and discussing its properties and role in assessing a decision-maker's preference. We present *game theory* and describe the decision-making setting we are investigating. The issue of ignorance in game theory is then highlighted, along with present solutions and the need for a probabilistic approach. Finally, we offer a summary of the following chapters.

1.1 Utility Theory

One of the central and main requirements of decision theories is knowing a decision-maker's preferences. Those preferences are reflected and quantified as *utilities*. As Von Neumann and Morgenstern [85] outlined, utilities are numerical measurable quantities that are expected to be complete. Furthermore, based on economics by Pareto and Bonnet [64], these utilities must be comparable. The decision-maker is expected to have a clear intuition of preference between two objects, events, or even a combination of events.

A utility function can measure a decision-maker's preferences over outcomes.

As suggested by Bernoulli [12], it is a non-linear function that allows determining the value of a reward. Formally, a utility function U will have a set of outcomes as domain and a set of real numbers \mathbb{R} as co-domain. Furthermore, it will be in line with the decision-maker's preferences; that is, if a decision-maker strictly prefers outcome ϕ_1 over ϕ_2 , then $U(\phi_1) > U(\phi_2)$.

In this study, to avoid misinterpretation, interest lies in the cardinal utility function that returns the 'worth' of each given outcome, in contrast to the ordinal utility function, which ranks outcomes based on preferences. The latter is a primitive concept often used in Economics and does not measure the degree of preference.

1.2 Problem Description

Game theory studies the decision-making process in a strategic setting, a situation where two or more decision-makers, known as players, compete or collaborate to maximise their utilities. Game theory relies on pre-determined utilities to decide a player's optimal move. In literature and practice, a player is required to know their utility towards each possible outcome in the game and the utilities of all other players [47]. Therefore, a player is expected to have explicit knowledge about each outcome, including outcomes that have never been previously experienced. This knowledge is reflected by having a utility function with the game's complete set of outcomes as a domain.

Usually, the utility function is considered to have a fixed form throughout the entire gameplay. Furthermore, it can be either a defined function, *e.g.* an exponential utility, or utility values attributed to each possible outcome. Whichever

way, the possibility that a player may not know in advance an outcome's payoff is often ignored, and a utility value is assumed. Hence, classical game-theoretic approaches fail to handle situations of uncertainty. Moreover, since an assumed utility value remains fixed throughout the entire gameplay, these approaches don't allow the player to be surprised or learn about their preferences once the outcome is experienced.

Given game theory's many applications in economy, finance, biology, law, marketing, management, etc., a state of uncertainty over outcomes could arise in many cases, *e.g.* when two companies invest in a new market. It is common for different parties to freely enter a game without having experienced its outcomes before, sometimes as a trial attempt to learn about their preferences towards these outcomes. Therefore, stretching the classical theory to allow for cases of ignorance over an outcome will expand its application. Work done by Nau [56] on extending de Finetti's [20] subjective probability approach to non-cooperative games, coupled with the imprecise probabilities toolkit, provides a compelling framework that could serve our purpose. It allows us to convert a game to a common-knowledge matrix and subsequently replace an unknown outcome's payoff with a gamble. This gamble will generate a random reward whose value depends on the result of an experiment. Inspired by Walley [86], an elicitation model that uses a player's pre-existing beliefs towards the gamble's outcomes is utilised to determine its supremum buying and infimum selling prices. Furthermore, Houlding and Coolen's [38] non-parametric predictive utility inference (NPUI) framework is used as a dynamic updating mechanism for repeated games. It allows players to change their preferences towards an unknown outcome based on new observations and therefore make more desirable moves.

1.3 Outline of Thesis

In the remainder of this report, Chapter 2 covers game theory's literature and applications in different strategic settings. Chapter 3 gives an overview of various approaches used to quantify, model, and assess uncertainty, e.g. imprecise probability, non-parametric utility inference (NPUI), etc. These approaches are our building blocks to a game theory model that allows for ignorance. Chapter 4 highlights the limitations of game theory under ignorance and suggests a framework that overcomes these limitations. The proposed solution covers one-shot and repetitive non-cooperative games. It includes an imprecise probability elicitation approach that allows players to provide initial utility previsions over uncertain outcomes and a dynamic updating model to adjust these elicited previsions whenever the relevant outcomes are experienced. In Chapter 5, extreme scenarios are considered, and the proposed framework is tested using a Monte Carlo simulation on a variant of the game 'Matching Pennies' that involves complete ignorance. Furthermore, this chapter explores how game-theoretic characteristics such as *zero-sum* and *symmetry* are altered by ignorance and proposes an approach to force these properties. Chapter 6 examines the implication of ignorance on utility diagnostics and introduces a model to improve the value of information (VOI) under this case. Finally, we conclude with Chapter 7 and propose possible developments to the proposed framework.

Review of Game Theory

This chapter covers the fundamental literature on game theory and its application in strategic settings. Section 2.1 briefly overviews game theory and formally defines a game and its components. Section 2.2 highlights the importance of information in games. Furthermore, it introduces *Value of Information* (VOI), an analytical method to quantify the value of a piece of information. Section 2.3 formally defines a game and the form used to represent it throughout this study. Section 2.4 discusses the different risk attitudes and their impact on a subject's utility function. Furthermore, it gives an introductory example of a game. Section 2.5 defines rationality and discusses the problem of rational behaviour within strategic settings. Finally, the chapter concludes with a discussion in Section 2.6 on game strategies and types of equilibria. We demonstrate through examples how to achieve each type of equilibrium.

2.1 Game Theory

Game theory studies the decision-making process in a competitive or collaborative setting. It aims at helping a rational and intelligent decision-maker (DM) select

the optimal choice amongst a set of available alternatives [52]. A DM is considered rational if they act in accordance with an accepted system of rationality axioms, which is discussed in Section 2.5. Furthermore, a DM is considered intelligent if they completely understand the situation under which a decision should be made and can make inferences that any other DM can make.

Game theory is used in different disciplines, including but not limited to economy, finance, biology, law, marketing, management, etc. The decision-making setting always involves more than one DM. We refer to them as players. They can either cooperate, *e.g.* two firms collaborating on a project, or compete, *e.g.* two firms fighting on market shares. In both cases, players are *interdependent*. The behaviour of one affects the other, either positively or negatively. Situations of interdependence are known as *strategic settings*. Under a strategic setting, a player must consider the actions of other players before making any decision. For example, if a firm is looking to maximise its shares in the market, it should analyse the position and activity of its rivals.

Although studies and research published around strategic settings are relatively new, logical thinking about human behaviour existed long ago. For instance, Parlor games, *e.g.* chess, were studied by mathematicians hundreds of years ago in an attempt to find optimal strategies. However, the first formal literature around games, made by Zermelo, appeared in 1913 [93]. Then, a ground-breaking work by Borel on the concept of strategies followed [14, 15]. In 1928 [84] and 1944 [85], Von Neumann and Morgenstern, in their revolutionary work, came up with a general theory for strategic situations called the *theory of games*. They meticulously outlined how games should be represented mathematically and proposed a general method to analyse them. However, it was limited to specific strategic settings.

Later on, in the mid-century, Nash [54] extended the method and distinguished between *cooperative* and *non-cooperative* games and created the concept of rational behaviour. After that, mathematicians and economists continued building on the established foundation and made the most powerful toolboxes used today in social sciences [85].

Developing a systematic understanding of a strategic setting and translating it into a quantitative model has several advantages. First, it provides a language to decrypt human behaviour under different strategic settings. Second, it allows assessing optimal behaviour within a strategic setting. Finally, it can be used as a framework to build other strategic settings.

Game theory requires a player to assign utilities to each possible outcome in the game. These utility values are used to assess the optimal choice available to this player and identify the strategy that maximises their expected utility. Since the game's outcome is a multilateral decision, when each player makes a choice, they should consider that other players are also trying to maximise their expected utility [47].

Games are classified into two main categories: *zero-sum* games, where the net total of all payments made and received by players is zero. In this case, the exchanged assets, *e.g.*, money, between players are not produced or destroyed during the game. *Non-zero-sum* games, where the total of payments is neither zero nor even constant [85].

Game theory is the standard approach to reasoning within strategic settings [10]. However, other approaches exist and are beyond the scope of this research. For instance, adversarial risk analysis (ARA) is an alternative to game theory. As highlighted by Banks et al. [10], it is a framework that guides a player's

decision-making process by analysing the game from their opponent's perspective. Uncertainties related to the opponent are partitioned into three separate components and then accounted for through subjective probability distributions [50]. ARA could lead to complex mathematical formulations. Nevertheless, it has a natural and intuitive concept. For example, consider a mental model of what an employer values (*e.g.* good performance, creativity, etc.) and their likely reactions to different discussions about a raise. The more accurate the model, the higher the likelihood of getting a raise.

2.1.1 Defining a Game

Prior to discussing the quantitative aspect of game theory, the following is an overview of a *game* and its components [85].

1. A game is the set of rules that describe it, and every instance of it is a play. For example, in the game Rock-Scissors-Paper, the rules state that each player can choose either rock, scissors or paper and that rock overpowers scissors, paper overpowers rock, scissors overpowers paper.
2. Moves are the component elements of a game. A move is an occasion where a player is presented with a set of alternatives to choose from, subject to the game's rules. Moves fall into two categories: a *personal* move, when a player's choice relies strictly on their free will, *e.g.* choosing scissors in the game Rock-Scissors-Paper. A *chance* move, when the player's choice is influenced by probability calculations, *e.g.* getting dealt a card;
3. Each player is free to base their decisions on pre-determined principles, as

long as these principles comply with the game's rules. The set of a player's principles is referred to as a *strategy* and is discussed further in Section 2.6.

2.2 Information in a Game

As previously mentioned, each player has to choose among several alternatives. If the player is at a *chance* move, they will rely solely on probability, and nobody's opinion influences their decision. However, suppose they are at a *personal* move. In that case, the player's state of information plays a vital role in making decisions because the available information helps them assess optimal decisions that maximise their expected utility.

At each *personal* move, the basic set of information accessible to each player should include all numerical conclusions derived from the game's rules; that is, the payoffs of all possible outcomes implied from the players' alternatives. Therefore, there is an assumption that the player knows the payoff of each outcome and has visibility of all the game's components [85].

Throughout the gameplay, a player becomes aware of all or some of the choices made by other players at their respective moves. A move is *anterior* to this player if it happens chronologically before their current move. However, it is said to be *preliminary* to the player only if they know the move's resulting decision. Therefore, anteriority doesn't necessarily imply preliminary. Anterior moves preliminary to a player will help them adjust their information set and reduce it by excluding the no longer possible payoffs.

Furthermore, the set of information is sensitive to any extra information regarding the game's components, *i.e.* possible outcomes and their payoffs. For instance,

if, through additional information, a player knows in advance what choices other players will make, they can narrow the range of possible payoffs further.

Value of information (VOI) is a method used in decision analysis to help a DM assess the value of an additional piece of information. VOI places a quantitative measure on available information representing the most a person would pay for this information. In the classic *Theory of Information*, Blackwell [13] assumes that the DM has a prior probability distribution over the possible states of the world. Before making a decision, they may be able to update their beliefs by observing available information, yielding a posterior probability distribution. Then the VOI is computed by calculating the difference between the expected return of the optimal decision based on prior beliefs and the one based on posterior beliefs. Therefore, when the amount of information is limited, or no information is available, VOI tends to zero [69].

2.3 Strategic Form of Games

The unilateral decision-making setting is a basic model that includes only one DM. The DM maximises their expected utility by choosing from a set of alternatives and their respective utilities. The *strategic* form of games extends this unilateral model to cover strategic settings that include more than one DM, *i.e.* players [59]. Furthermore, it makes it possible to determine the entire course of a game-play through each player's plan of action, known as *strategy*. The *strategic* form helps bring the formal definition of games to the following description [85], which is used in the remainder of this study.

1. Only *personal* moves matter;

2. Each player $k = 1, \dots, K$ has a finite set of alternatives $\Lambda^k = \{a_1^k, \dots, a_{m^k}^k\}$, where K is the total number of players, and m^k is the total number of alternatives available to player k ;
3. When players choose an alternative from their respective set Λ^k , they are not aware of each other's decisions;
4. The matrix $\Phi = \Lambda^1 \times \Lambda^2 \times \dots \times \Lambda^K$, composed of K -dimensional vectors, represents all possible combinations of alternatives across players; that is, the game's set of possible outcomes. Therefore, after players make their decisions, the outcome of the play is $\phi \in \Phi$;
5. For any outcome $\phi \in \Phi$, each player k pays or receives a payment determined by the payoff function $r^k(\cdot)$. For example, outcome ϕ has a payoff $r^k(\phi)$. This payoff is denominated in an arbitrary single monetary commodity. Furthermore, it is considered unrestrictedly divisible, substitutable, and freely transferable;
6. Each player k has a utility function $U^k(\cdot)$ whose domain is the set of possible payoffs in the game. Hence, player k 's utility towards outcome $\phi \in \Phi$ is $U^k(r^k(\phi))$.

Although the *strategic* form supports games with sequential decision-making, a more natural model for these cases is the *extensive* form introduced by Von Neumann and Morgenstern [85] and Kuhn [43]. In games with *perfect* information, the *extensive* form allows each player to know the previous decisions made by others. However, using the *strategic* form suffices for this study. Hence, the *extensive* form is out of scope.

2.4 Risk Attitudes

Besides its role in measuring a player's preferences over outcomes, a utility function U^k reflects a player's risk attitude towards the payoff of these outcomes [40]. This risk attitude translates to the level of risk a player k is willing to accept for a particular payoff $r^k(\phi)$. Furthermore, it directly influences the curvature and shape of the utility function.

Consider a gamble where a DM has equal chances to win 5,000\$ or lose 1,000\$. The fair actuary premium of this gamble is 3,000. However, the amount that the DM is willing to pay depends on their risk attitude. As Arrow [1] and Pratt [68] discussed, a DM who prefers to avoid actuarially fair gambles is considered risk-averse. Conversely, a DM who prefers pursuing these gambles is considered risk-seeking. However, a DM indifferent about actuarially fair gambles is deemed risk-neutral.

In practice, the most common risk attitudes are risk-neutral and risk-averse. Risk-neutral decision-making is often used for situations involving small risks, whereas risk-averse decision-making is more frequent for situations with considerable risks. Although a risk-seeking attitude is possible in real life, *e.g.* in entrepreneurs, this study is limited to the risk-neutral and risk-averse cases [41].

In game theory, when a player is risk-neutral, their choices are not affected by the underlying risk of the available outcomes. For instance, if two outcomes generate the same payoffs, however, with different levels of a relevant risk, *e.g.* , losing a premium, the player is indifferent about which outcome to choose.

A risk-neutral player's utility towards an outcome $\phi \in \Phi$ is the reward of that outcome, *i.e.* its payoff $r^k(\phi)$. This utility is a linear function of the form

$U^k(r^k(\phi)) = a[r^k(\phi)] + b$, where $a > 0$, and $b \in \mathbb{R}$. Since a utility function is unique up to a positive affine transformation, the payoff function r^k can be used to assess the utility of each outcome.

As opposed to the risk-neutral case, if a player is risk-averse, they prefer the outcome that returns the optimal reward for a particular level of risk. Therefore, between two outcomes with similar payoffs, they choose the one with the lowest level of risk.

A risk-averse player's utility function U^k is concave downward over the set of payoffs r^k [68, Theorem 1]. In the remainder of this study, we use an exponential utility function to assess a risk-averse player's utility. However, that's not a restriction, and any relevant concave function can be used instead. This exponential utility function is formally defined as follows:

$$U^k(r^k(\phi)) = 1 - \exp(-c \times r^k(\phi)), \quad (2.1)$$

where $c > 0$ represents a risk aversion parameter.

Example 2.4.1. *Consider the classic game 'battle of the sexes', introduced by Luce and Raïffa [47], where two players plan to go for entertainment. They prefer going together rather than by themselves. Furthermore, if they go together, the player going to their favourite place gets more satisfaction than the other.*

Assume that players are risk-neutral and have two alternatives, either 'Hockey' or 'Cinema'. Player one prefers going to 'Hockey', whereas player two prefers going to the Cinema'.

This game involves two players ($z = 2$). Therefore, converting it to its strategic form is done through a $m^1 \times m^2$ payoff matrix representing the provided informa-

tion. As seen in Table 2.1, player one has the following $m^1 = 2$ alternatives:

- Hockey, denoted by Top;
- Cinema, denoted by Bottom.

Player two also has the following $m^2 = 2$ alternatives:

- Hockey, denoted by Left;
- Cinema, denoted by Right.

Hence, $\Lambda^1 = \{\text{Top}(T), \text{Bottom}(B)\}$, $\Lambda^2 = \{\text{Left}(L), \text{Right}(R)\}$, and all possible outcomes in the game are $\Phi = \Lambda^1 \times \Lambda^2 = \{TL, TR, BL, BR\}$.

Each cell in Table 2.1 shows both players' payoffs. For example, for outcome TL , i.e. row Top and column Left, the values $(2,1)$ indicate that if players one and two go together to a hockey game, their respective payoffs would be $r^1(TL) = 2$ and $r^2(TL) = 1$. Conversely, for outcome TR , if these players go separate ways, player one to a hockey game and two to the cinema, their payoffs would be $r^1(TR) = r^2(TR) = 0$.

Table 2.1: 'Battle of the sexes' - Payoff matrix.

	Left	Right
Top	2,1	0,0
Bottom	0,0	1,2

If players are risk-neutral, their utility towards an outcome ϕ is $U^k(r^k(\phi)) = r^k(\phi)$. Hence, the payoff matrix in Table 2.1 can also represent these players' utility towards each outcome. However, as seen in Example 3.2.2, if players are risk-averse, the utility function in Equation (2.1) is used to denominate players' payoffs in utility currency.

2.5 Rational Behaviour

Like an economic community that involves gaining and spending incomes, the analysis of a game relies on understanding the behaviour of its agents, *i.e.* the players. Amongst the many difficulties that behavioural analysis entails, correctly describing the players' assumed motives is one of the main obstacles. However, this obstacle is overcome by the assumption that all players seek to maximise a single monetary commodity, an asset freely transferable, divisible, and substitutable with utility. Therefore, a player that aims to maximise utility is considered rational. In game theory, rationality is formally defined as follows: let $\phi_i^k \in \Phi$ denote an outcome in which a player k chooses the alternative a_i^k . Let $\Lambda_i^k \subseteq \Lambda^k$ denote the subset of all outcomes that include the alternative a_i^k . Player k is rational if and only if they choose the alternative a_i^k that solves the problem $\max_{a_i^k \in \Lambda_i^k} U^k(r^k(\phi_i^k))$, or $\max_{a_i^k \in \Lambda_i^k} r^k(\phi_i^k)$ if the player k is risk-neutral [61].

A strategic setting involves multiple players who, if rational, each seek to optimise their utility. Therefore, a game's outcome is not only influenced by the decision of a particular player but is the result of a combination of decisions made by all players. Moreover, each player has several alternatives and possibly different information and understanding of the game. Hence, they have multiple ways to maximise their utility.

Since a player has several courses of action, rational behaviour is a complete set of rules that state how they should behave in any possible situation. These rules are called *strategies*, and each strategy should indicate the minimum utility a player gets if they act rationally or if other players behave irrationally.

When a strategy brings stability and balance to a game, *i.e.* an equilibrium

where all players achieve an optimal utility, it becomes a solution to that game and the problem of rational behaviour.

2.6 The Strategy

During a game-play, each player adopts a strategy that can be either *pure* or *mixed*. A *pure* strategy is a complete plan that helps the player make a choice from a set of alternatives for every conceivable scenario.

A *mixed* strategy is a probability distribution over pure strategies; it allows randomising pure strategies. It provides the player with a vector of probabilities that includes the probability of playing each alternative. A mixed strategy is beneficial for situations where discovering the opponent's decision is what matters most. Furthermore, it aids in concealing the player's intention by pushing them to play the available alternatives irregularly [85].

A Nash equilibrium is a state where no player can gain a higher payoff by unilaterally deviating from their strategy. If no Nash equilibria exist in pure strategies, then at least one must exist in mixed strategies where each player makes choices based on a probability distribution over their set of alternatives [53].

In 1974, Aumann studied the impact of players correlating their choices [6]. He used a randomisation device on outcomes where players may disagree. This approach revealed that a *correlated* mixed strategy could lead to strictly higher expected payoffs than Nash equilibria. Furthermore, it could remove the competitive aspect from non-cooperative games and push players to cooperate.

The following sections discuss Nash and correlated equilibria with further details, showing how to achieve each type of equilibrium. For the sake of simplicity,

the provided examples only consider the case of risk-neutral players.

2.6.1 Pure Nash Equilibria

A pure Nash equilibrium is an outcome reached through a tuple of pure strategies, one strategy per player, where no player can gain a higher utility by solely deviating from their strategy [53]. It is only possible for a player to get a higher utility if other players change their strategies. In other words, an outcome is a pure Nash equilibrium if and only if all players play their best responses to each other's pure strategies; for a given state, each player chooses the alternative that returns the highest utility.

A game can have several pure Nash equilibria. For example, in Table 2.1, both outcomes TL and BR are pure Nash equilibria. Players one and two cannot gain more utility if each unilaterally deviates from their pure strategy. Therefore, at equilibrium TL , if player one changes their strategy and chooses the alternative B , the game's outcome becomes BL , resulting in zero utility for both players.

The game in Example 2.4.1 is symmetric, an intuitive concept referring to a case where a game is the same for all players. Hence, if both players switch places in 'battle of the sexes', their payoffs remain the same. Therefore, similar to player one, if player two deviates from equilibrium TL , the game's outcome becomes TR , resulting in zero utility for both players.

If a Nash equilibrium does not exist in pure strategies, one must exist in mixed strategies [53]. In this case, the Nash equilibrium is reached when each player's mixed strategy is their best response.

2.6.2 Mixed Nash Equilibria

A mixed Nash equilibrium is a profile of mixed strategies, one strategy per player, where none of these players can get a higher expected utility by solely deviating from their mixed strategy. Each player's mixed strategy represents their best response to the other players' mixed strategies. Therefore, the mixed Nash equilibrium leads to a steady state where each player randomises their choice of alternatives according to an unchanged vector of probabilities. These randomisations are statistically independent between players.

Let $\vec{\varepsilon}_k$ denote a vector of probabilities representing player k 's mixed strategy. Let ε^* denote a profile of mixed strategies. ε^* is a mixed Nash equilibrium if and only if, for each player k , the expected utility generated by ε^* is at least as great as the expected utility of every other mixed strategy profile $(\vec{\varepsilon}_k, \vec{\varepsilon}_{-k})$, where $\vec{\varepsilon}_{-k}$ is a notation used to represent the list of mixed strategies played by players other than k . Formally,

$$U^k(\varepsilon^*) \geq U^k(\vec{\varepsilon}_k, \vec{\varepsilon}_{-k}), \quad (2.2)$$

for all possible mixed strategies $\vec{\varepsilon}_k$ of each player k , where $U^k(\varepsilon^*)$ is player k 's expected utility for a mixed strategy profile ε^* .

In a two-player game with a mixed strategy profile $(\vec{\varepsilon}_1, \vec{\varepsilon}_2)$, player k 's expected game utility is computed as follows:

$$\begin{aligned} U^k(\vec{\varepsilon}_1, \vec{\varepsilon}_2) &= \sum_{\phi \in \Phi} E^k [U^k(r^k(\phi))] \\ &= \sum_{a_i^1 \in \Lambda^1} \sum_{a_j^2 \in \Lambda^2} U^k(r^k(a_i^1 a_j^2)) \varepsilon_{a_i^1} \varepsilon_{a_j^2}, \end{aligned} \quad (2.3)$$

where $\varepsilon_{a_i^k}$ denotes the probability of player k choosing alternative $a_i^k \in \Lambda^k$.

At a mixed Nash equilibrium, each player k 's vector of probabilities $\vec{\varepsilon}_k$ makes the opposite players indifferent to which alternatives to choose. For each player, all available alternatives return the same expected utility. For example, in a two-player, two-alternative game, the profile $(\vec{\varepsilon}_1, \vec{\varepsilon}_2)$ is a Nash equilibrium if and only if:

- Player two's expected utility $\sum_{a_j^1 \in \Lambda^1} U^2(r^2(a_i^1 a_j^2)) \varepsilon_{a_j^1}$ is the same for all alternatives $a_j^2 \in \Lambda^2$;
- Player one's expected utility $\sum_{a_j^2 \in \Lambda^2} U^1(r^1(a_i^1 a_j^2)) \varepsilon_{a_j^2}$ is the same for all alternatives $a_i^1 \in \Lambda^1$.

Example 2.6.1. *Cont'd.* In the game 'battle of the sexes', assume that both players are risk-neutral. Hence, $U^1(r^1(\phi)) = r^1(\phi)$ and $U^2(r^2(\phi)) = r^2(\phi)$. Given player one's strategy $\vec{\varepsilon}_1$, player two's expected utility for alternative L is:

$$r^2(TL)\varepsilon_T + r^2(BL)\varepsilon_B,$$

and their expected utility for alternative R is:

$$r^2(TR)\varepsilon_T + r^2(BR)\varepsilon_B.$$

Player two is indifferent between alternatives L and R if they get the same expected utility for any for them. Formally, that's when:

$$r^2(TL)\varepsilon_T + r^2(BL)\varepsilon_B = r^2(TR)\varepsilon_T + r^2(BR)\varepsilon_B.$$

Since $\varepsilon_T + \varepsilon_B = 1$, alternatives L and R return the same expected utility when:

$$\varepsilon_T = \frac{r^2(BR) - r^2(BL)}{r^2(TL) - r^2(BL) - r^2(TR) + r^2(BR)} = \frac{2}{3}$$

$$\varepsilon_B = \frac{r^2(TL) - r^2(TR)}{r^2(TL) - r^2(BL) - r^2(TR) + r^2(BR)} = \frac{1}{3}.$$

Similarly, given player two's strategy $\vec{\varepsilon}_2$, player one is indifferent between alternatives T and B if they get the same expected utility for any of them. Formally, that's when $\vec{\varepsilon}_2 = \{\varepsilon_L = \frac{1}{3}, \varepsilon_R = \frac{2}{3}\}$.

A Nash equilibrium is when players one and two respectively play the strategies $\vec{\varepsilon}_1$ and $\vec{\varepsilon}_2$. Therefore, the profile $\varepsilon^* = (\vec{\varepsilon}_1, \vec{\varepsilon}_2)$ is a mixed Nash equilibrium. Using Equation 2.3, the expected utility of profile ε^* is $\frac{2}{3}$ for both players.

2.6.3 Correlated Equilibria

In 1987, Aumann defined a correlated equilibrium as a function that maps a finite probability space to the set of all possible outcomes Φ [8]. He considered a Nash equilibrium to be a particular case of it. Unlike Nash equilibria, one of the most interesting aspects of correlated equilibria is the ease of computing them by simply solving a system of linear inequalities.

Let $\vec{\rho}$ denote a correlated strategy. It is a probability vector in which an element ρ_ϕ is the probability of the outcome $\phi \in \Phi$. In a two-person game, $\vec{\rho}$ is a correlated

equilibrium if and only if:

$$\left\{ \begin{array}{l} \rho_{\phi} \geq 0, \forall \phi \in \Phi, \\ \vec{\rho}' \mathbf{1} = 1, \\ \sum_{a_j^2 \in \Lambda^2} (r^1(a_i^1 a_j^2) - r^1(a_q^1 a_j^2)) \rho_{a_i^1 a_j^2} \geq 0, \forall a_i^1, a_q^1 \in \Lambda^1, \text{ and } i \neq q, \\ \sum_{a_i^1 \in \Lambda^1} (r^2(a_i^1 a_j^2) - r^2(a_i^1 a_k^2)) \rho_{a_i^1 a_j^2} \geq 0, \forall a_j^2, a_k^2 \in \Lambda^2, \text{ and } j \neq k. \end{array} \right. \quad (2.4)$$

Note that $\vec{\rho}'$ is the transpose of $\vec{\rho}$.

The system of linear inequalities (2.4) defines a convex polytope of correlated equilibria. This polytope is the bounded intersection of a finite set of closed half-spaces and contains all the rational solutions of the game.

For a correlated strategy $\vec{\rho}$, each player k 's expected utility is computed as follows:

$$\begin{aligned} U^k(\vec{\rho}) &= \sum_{\phi \in \Phi} E^k [U^k(r^k(\phi))] \\ &= \sum_{\phi \in \Phi} r^k(\phi) \rho_{\phi}. \end{aligned} \quad (2.5)$$

Let N denote the number of possible outcomes in Φ . In System (2.4), the two constraints $\vec{\rho} \geq 0$ and $\vec{\rho}' \mathbf{1} = 1$ define a $N - 1$ dimensional simplex, containing all probability distributions on outcomes Φ . The polytope defined by all constraints is a subset of this simplex. If the polytope has a dimension smaller than $N - 1$, the distribution of correlated equilibria will lie on its boundary. Let I denote the set of all joint probability distributions of independent variables (here, players). It

is the system of nonlinear constraints,

$$I = \{ \vec{\rho} : \rho_{\phi} = \rho_1(a^1) \times \rho_k(a^k) \times \dots \times \rho_z(a^z), \forall \phi \in \Phi \},$$

where $a^k \in \Lambda^k$ and ρ_k is the marginal probability distribution on Λ^k . In a 2×2 game, Nau et al. [57] describe the resulting simplex as a 3-dimensional tetrahedron and I as a 2-dimensional saddle. The set of Nash equilibria is the intersection of I and the correlated equilibria polytope. Nash equilibria only rest on the surface of this polytope.

Example 2.6.2. *Cont'd.* In ‘battle of the sexes’ with risk-neutral players, the system of linear inequalities (2.4) results in the following polytope, denoted by C .

$$C = \begin{cases} 2\rho_{TL} - \rho_{TR} \geq 0, \\ -2\rho_{BL} + \rho_{BR} \geq 0, \\ \rho_{TL} - 2\rho_{BL} \geq 0, \\ -\rho_{TR} + 2\rho_{BR} \geq 0, \\ \rho_{\phi} \geq 0, \forall \phi \in \{TL, TR, BL, BR\}, \\ \rho_{TL} + \rho_{TR} + \rho_{BL} + \rho_{BR} = 1. \end{cases}$$

Figure 2-1 is a plot of polytope C , generated using the python code [28]. As seen, C is a hexahedron with five vertices, listed in Table 2.2. Furthermore, Figure 2-1 shows three intersections between the polytope C and the set of all joint probability distributions. Intersections TL and BR are pure Nash equilibria, whereas the intersection $(2/3 T, 1/3 B) \times (1/3 L, 2/3 R)$ is a mixed Nash equilibrium that sits





Figure 2-1: 'Battle of the sexes' polytope and set of all joint independent probability distributions between players - Intersections TL and BR are pure Nash equilibria, and intersection $(2/3 T, 1/3 B) \times (1/3 L, 2/3 R)$ is a mixed Nash equilibrium.

at the inefficient frontier of the polytope, hence it is not a Pareto efficient strategy.

Table 2.2: ‘Battle of the sexes’ with risk-neutral players - Polytope vertices.

	TL	TR	BL	BR
Vertex 1	1	0	0	0
Vertex 2	0	0	0	1
Vertex 3	2/9	4/9	1/9	2/9
Vertex 4	2/5	0	1/5	2/5
Vertex 5	1/4	1/2	0	1/4

Imprecise Probabilities and Related Literature

This chapter briefly overviews different theories used to quantify, model, and assess uncertainty. These theories are building blocks to the model proposed in subsequent chapters. Section 3.1 introduces the *imprecise probabilities* framework. Section 3.2 discusses how a game can be transformed into a revealed-rules matrix using de Finetti's approach to defining subjective probabilities, which utilises bets to assess beliefs. It also discusses how a revealed-rules matrix can compute the correlated equilibria of a game. Section 3.3 presents a non-parametric predictive utility inference framework and a utility learning model that, under a non-strategic setting, helps a decision-maker learn about their preference towards a novel outcome once it is experienced.

3.1 Imprecise Probabilities

Imprecise probabilities is a well-established framework aimed toward quantification and inference under uncertainty, a state of incomplete or vague information.

Influenced by de Finetti's work [21] on subjective probability, Williams [88, 89] worked on an early detailed study of the theory, and Walley [86] then developed it further to a more mature one.

The *imprecise probabilities* framework goes beyond classic approaches based on precise probabilities, where an event is assigned a single probability. It extends the traditional theories by allowing for incompleteness, imprecision, and indecision. Therefore, it permits modelling situations where relying on conclusions from incomplete information is important.

3.1.1 Lower and Upper Previsions

An attractive theory under imprecise probabilities is that of lower and upper previsions, represented by $\underline{P}(\cdot)$ and $\overline{P}(\cdot)$. They are, respectively, the supremum acceptable buying price and the infimum acceptable selling price of a gamble. Let \mathcal{X} denote a finite set of an experiment's exhaustive and mutually exclusive outcomes.

Let $f(\cdot)$ denote an arbitrary reward function defined on \mathcal{X} : if x is the outcome of the experiment, then the reward is $f(x)$, denominated in units of a linear utility-scale. As the experiment's outcome is random, the reward of the experiment is random. Hence, the reward of the experiment is interpreted as a gamble.

The lower and upper previsions of f are, respectively, $\underline{P}(f)$ and $\overline{P}(f)$. Furthermore, they are considered a subject's lower and upper expectations of f . Whenever they coincide, such that $P(f) := \underline{P}(f) = \overline{P}(f)$, the resulting $P(f)$ is called a linear prevision. De Finetti sees it as the fair price of f [21].

3.1.2 Desirability

Consider a subject who is publicly willing to buy a gamble f for a supremum acceptable price $\underline{P}(f)$, in arbitrary small positive quantities α . This decision indicates that they are willing to accept a payoff vector of $\alpha(f(\mathcal{X}) - \underline{P}(f))$. For example, let $\mathcal{X} = (O_1, O_2, O_3)$ and $f(\mathcal{X}) = (0, 3, 5)$. Buying f for $\underline{P}(f) = 3$ is equivalent to accepting the payoff vector $(-3\alpha, 0\alpha, 2\alpha)$.

A gamble is said to be *desirable* when a subject is willing to accept it whenever offered. This doesn't necessarily suggest that a non-desirable gamble is rejected. A non-desirable gamble only reflects that the subject is undecided about whether to accept it or not. Formally, in Walley's axioms of desirability [86, p.60], axiom (D1) considers a gamble f to be desirable if $\inf f > 0$, *i.e.* when it increases the subject's utility no matter the outcome.

A gamble f is said to be *almost desirable* if $\forall \epsilon > 0$, $f + \epsilon$ is desirable. The set of almost-desirable gambles includes all desirable gambles. Furthermore, as per [86, Theorem 3.8.4, p.158], these almost-desirable gambles correspond to some linear previsions that are greater than or equal to zero: if f is almost desirable, then $P(f) \geq 0$. This condition and the set of almost-desirable gambles are useful for eliciting lower previsions, which is demonstrated in the subsequent chapters.

3.1.3 Conditional Previsions

Let \mathcal{A} denote a subset of \mathcal{X} and $\delta_{\mathcal{A}}$ a 0-1 binary vector indicating whether the experiment's outcome belongs to \mathcal{A} (1 if true). Continuing the example in Section 3.1.2, let $\mathcal{A} = \{O_2, O_3\}$, hence, $\delta_{\mathcal{A}} = (0, 1, 1)$. A conditional lower prevision $\underline{P}(f|\mathcal{A})$ is the supremum buying price of gamble f given \mathcal{A} , such that $(f(\mathcal{X}) - \underline{P}(f|\mathcal{A}))\delta_{\mathcal{A}}$ is

desirable. Here, $\underline{P}(f|\mathcal{A})$ is a scalar, and the multiplication is performed element-wise.

A subject who is publicly willing to buy f for a supremum price $\underline{P}(f|\mathcal{A})$, in arbitrary small positive quantities α , is ready to accept a payoff vector of $\alpha(f(\mathcal{X}) - \underline{P}(f|\mathcal{A}))\delta_{\mathcal{A}}$. Hence, buying f given \mathcal{A} for $\underline{P}(f|\mathcal{A}) = 2$ is equivalent to accepting the payoff vector $(0, \alpha, 3\alpha)$.

A conditional lower prevision has two different interpretations. $\underline{P}(f|\mathcal{A})$ could be considered the supremum buying price of a gamble f whose value is zero outside \mathcal{A} . Or, as Walley's [86] *updating principle* suggests, it could be considered as the updated supremum buying price of f after receiving information that the outcome belonged to \mathcal{A} .

3.1.4 Lower and Upper Probabilities

A particular case of upper and lower previsions is when the reward function takes the form of an indicator function, *e.g.* $f(\mathcal{X}) = \delta_{\mathcal{A}}$. In that case, these previsions can be interpreted as lower and upper probabilities of the event that the experiment's outcome belongs to \mathcal{A} : $\underline{P}(\mathcal{A}) := \underline{P}(f)$ and $\bar{P}(\mathcal{A}) := \bar{P}(f)$. The lower probability $\underline{P}(\mathcal{A})$ represents all evidence certainly in favour of \mathcal{A} , and the upper probability $\bar{P}(\mathcal{A})$ represents all evidence possibly in favour of \mathcal{A} .

3.1.5 Coherence

A lower or upper prevision assessment should not allow any opportunities for riskless profits through intelligent combinations. It is a rationality requirement known as *coherence*. On a linear space, Walley [86] characterizes *coherence* using

the following axioms:

- $\underline{P}(f) \geq \inf f$ and $\overline{P}(f) \leq \sup f$;
- $\underline{P}(\lambda f) = \lambda \underline{P}(f)$ and $\overline{P}(\lambda f) = \lambda \overline{P}(f)$, $\forall \lambda \geq 0$;
- $\underline{P}(f + g) \geq \underline{P}(f) + \underline{P}(g)$ and $\overline{P}(f + g) \leq \overline{P}(f) + \overline{P}(g)$.

Coherent lower or upper previsions are challenging to achieve, especially in a non-linear space. Therefore, the lower and upper previsions theory allows constructing coherent models from assessments that only avoid sure loss. *Avoiding sure loss* is a weaker rationality condition that is easier to satisfy. It requires rejecting any transaction that generates a sure loss no matter what the outcome of the gamble is; failing to do so is a fundamental irrationality [86, p.67].

3.2 Revealed-Rules Matrix

Nau [56] utilised imprecise probabilities to create a unified theory between subjective probability, game theory, and other equilibrium models used for games and markets. His approach allows converting a non-cooperative game into a revealed-rules matrix that reveals the rules required to generate a convex set of probability distributions representing the game's equilibria.

3.2.1 Risk Neutral Players

Consider a non-cooperative game with K risk-neutral players. Let Λ^k denote a finite set of alternatives available to player k , and $\Phi = \prod_{k \in K} \Lambda^k$ denote the set of all possible outcomes in the game. Let r^k denote a $|\Phi|$ -dimensional vector that

represents the payoff of player k as a function of these outcomes. Therefore, for outcome $\phi \in \Phi$, player k 's payoff is $r^k(\phi)$.

Consider e_i^k to be the event in which player k chooses the alternative $a_i^k \in \Lambda^k$ over any other alternatives. Let r_i^k denote the payoff vector available to player k after making this choice, *i.e.* $r_i^k = (r^k|e_i^k)$. For example, in the game paper-rock-scissors, if player one chooses to play 'rock', the resulting payoff vector r_i^k has the payout of outcomes *rock - rock*, *rock - scissors*, *rock - paper*.

The occurrence of event e_i^k means that player k would trade any payoff vector r_j^k ($j \neq i$) for r_i^k , conditional on e_i^k . This trade-off translates into an unconditional bet with a *true* payoff vector of $(r_i^k - r_j^k)e_i^k$. Conversely, suppose a player chooses to publicly accept any small bet whose payoff vector is proportional to $(r_i^k - r_j^k)e_i^k$. In that case, they are making their *true* payoff function common knowledge at the discretion of any observer [7]. Following the same logic, if all players in the game are willing to accept small conditional bets, they make their *true* payoff function public knowledge. As a result, a matrix that exhaustively lists each player's possible *true* payoff functions is built. This matrix is defined as follows:

Definition 1. *A game's revealed-rules(GRR) matrix, denoted by \mathbf{M} , represents the true payoff function of each player for each possible bet they could accept. Matrix \mathbf{M} 's columns are indexed by outcomes in Φ , and its rows are indexed by $r_i^k - r_j^k$. As described by [56] and demonstrated below, a GRR matrix contains all 'commonly-knowable information about the rules that is actually used in determining the equilibria of non-cooperative games.'*

Nau [58] shows that for risk-neutral players, *i.e.* players whose utility function is state-independent, to play the game rationally, they should have jointly coher-

ent strategies. This condition is fulfilled if and only if a correlated equilibrium exists. Given the correlated strategy $\vec{\rho}$ and the revealed-rules matrix \mathbf{M} , players are rational if and only if

$$\begin{cases} \vec{\rho} \geq 0, \\ \vec{\rho}'\mathbf{1} = 1, \\ \mathbf{M}\vec{\rho} \geq 0. \end{cases} \quad (3.1)$$

Example 3.2.1. *Cont'd.* Table 3.1 shows the resulting revealed-rules matrix, \mathbf{M} , of 'battle of the sexes'. For instance, the first row represents the case where player one chooses alternative T over B . This choice exposes them to two possible transactions, $r^1(TL) - r^1(BL) = 2$, and $r^1(TR) - r^1(BR) = -1$. Furthermore, System (3.1) returns a set of linear inequalities resulting in polytope C . It can be noted from Example 2.6.2 that these linear inequalities are the same as the linear inequalities produced by System (2.4), which defines correlated equilibria.

Table 3.1: 'Battle of the sexes' - Revealed-rules matrix \mathbf{M} .

	TL	TR	BL	BR
$r_T^1 - r_B^1$	2	-1	0	0
$r_B^1 - r_T^1$	0	0	-2	1
$r_L^2 - r_R^2$	1	0	-2	0
$r_R^2 - r_L^2$	0	-1	0	2

3.2.2 Risk Averse Players

Nau also examined the more generic case of risk-averse players who hedge and accept less risky transactions than the risk-neutral ones. The non-linear utility of these players is affected by the smallest transaction. Therefore, risk aversion introduces more imprecision and requires a model that considers state-dependent

marginal utility for money. Furthermore, as seen later, it leads to a more extensive set of correlated equilibria than the risk-neutral case.

Let U^k denote a $|\Phi|$ -dimensional vector that represents the utility of player k as a function of the payoffs of the game's outcomes. Therefore, for outcome $\phi \in \Phi$, player k 's utility is $U^k(r^k(\phi^k))$. Let \dot{U}^k denote the derivative of U^k with respect to the payoff $r^k(\phi^k)$, and σ^k a payoff value that represents risk. Suppose that σ^k is larger than the player's risk tolerance.

As opposed to the risk-neutral case where a player accepts a transaction according to its expected utility, a risk-averse player k agrees to a transaction based on its marginal expected utility with respect to σ^k . Specifically, this transaction is acceptable when adding its payoff to σ^k positively impacts the overall expected utility. Formally, if player k 's current expected utility is $E^k[U^k(\sigma^k)]$, transaction θ is acceptable if $E^k[U^k(\sigma^k + \theta)] \geq E^k[U(\sigma)]$. Alternatively, if θ is small in magnitude, θ is acceptable if $E^k[\dot{U}^k(\sigma^k)\theta] \geq 0$.

Nau extends the GRR matrix discussed in Section 3.2.1 to cover risk aversion. Let U_i^k denote the payoff vector available to player k for event e_i^k , *i.e.* $U_i^k = (U^k|e_i^k)$. The occurrence of event e_i^k means that risk-averse player k would trade any utility vector U_j^k ($j \neq i$) for U_i^k , conditional on e_i^k . This trade-off translates into an unconditional bet with a *true* utility vector of $(U_i^k - U_j^k)e_i^k$. Conversely, suppose a player chooses to publicly accept any small bet whose payoff vector is proportional to $(U_i^k - U_j^k)e_i^k$. In that case, they are making their *true* utility function common knowledge at the discretion of any observer [7], allowing the construction of a GRR matrix (refer to Section 3.2.1).

In the risk-averse case, each bet $U_i^k - U_j^k$ is denominated in *utility* units and should be converted to commonly understood units of money. This conversion is

possible by dividing $U_i^k - U_j^k$ by the relevant vector of local marginal utility of money, *i.e.* \dot{U}_i^k . The bet's resulting *true* utility vector denominated in monetary value becomes:

$$\frac{U_i^k - U_j^k}{\dot{U}_i^k}.$$

Similarly to the risk-neutral case, if all players in the game are willing to accept small conditional bets, they make their *true* utility function public knowledge. As a result, a revealed-rules matrix \mathbf{M}^* is built. It exhaustively lists each player's possible *true* utility functions. Matrix \mathbf{M}^* 's columns are indexed by outcomes in Φ and its rows are indexed by $\frac{U_i^k - U_j^k}{\dot{U}_i^k}$.

Nau [56] shows that risk-averse players have jointly coherent strategies if and only if a risk-neutral equilibrium exists; that is, if at least one correlated equilibrium exists. In this case, the set of correlated equilibria is computed by substituting \mathbf{M} with \mathbf{M}^* in the System (3.1), then solving the resulting inequalities.

Example 3.2.2. *Cont'd.* Consider that players in 'battle of the sexes' are risk-averse and have an exponential utility of the form

$$U^k(r^k(\phi)) = 1 - \exp(-LN(\sqrt{2}) \times r^k(\phi)),$$

where $LN(\sqrt{2})$ represents a particular risk aversion parameter. Table 3.2 shows this game's resulting payoff matrix, denominated in utility units. The derivative of

Table 3.2: 'Battle of the sexes' with risk-averse players - Payoff matrix, denominated in utility units.

	Left	Right
Top	0.5,0.293	0,0
Bottom	0,0	0.293,0.5

U^k is $\dot{U}^k(r^k(\phi)) = LN(\sqrt{2})\exp(-LN(\sqrt{2}) r^k(\phi))$, and the resulting revealed-rules matrix, denominated in monetary value, is as follows:

Table 3.3: ‘Battle of the sexes’ with risk-averse players - Revealed-rules matrix M^* .

	TL	TR	BL	BR
$\frac{U_T^1 - U_B^1}{\dot{U}_T^1}$	2.885	-0.845	0	0
$\frac{U_B^1 - U_T^1}{\dot{U}_B^1}$	0	0	-1.442	1.195
$\frac{U_L^2 - U_R^2}{\dot{U}_L^2}$	1.195	0	-1.442	0
$\frac{U_R^2 - U_L^2}{\dot{U}_R^2}$	0	-0.845	0	2.885

In System (3.1), substituting matrix \mathbf{M} with \mathbf{M}^* results in the polytope of Figure 3-1, generated using the python code [28]. Similar to Figure 2-1, this polytope is a hexahedron with five vertices listed in Table 3.4. However, it is greater in size than that of the risk-neutral case, indicating a more extensive set of equilibria.

Table 3.4: ‘Battle of the sexes’ with risk-averse players - Polytope vertices.

	TL	TR	BL	BR
Vertex 1	0.35	0	0.29	0.35
Vertex 2	1	0	0	0
Vertex 3	0.18	0.63	0	0.18
Vertex 4	0.16	0.54	0.13	0.16
Vertex 5	0	0	0	1





Figure 3-1: 'Battle of the sexes' with risk-averse players - Polytope and set of all joint independent probability distributions between players.

3.3 Non-parametric Utility Updating

Based on Augustin and Coolen’s work [4], Houlding and Coolen [38] introduced a non-parametric predictive utility inference (NPUI) framework for utility induction under extreme ignorance. Their work considers decision-making within non-strategic settings.

For sequential decision problems, NPUI featured an attractive updating mechanism that assesses the impact of additional observations on utility previsions. This updating mechanism allows the creation of a learning model that adjusts a decision-maker’s utility previsions towards a novel outcome whenever this outcome is experienced.

NPUI’s updating mechanism is based on Hill’s assumption $A_{(n)}$ [34, 35, 36]. $A_{(n)}$ is particularly useful for predictions with extremely vague prior knowledge of the underlying distribution. It assumes that pre-observations are exchangeable [21], that is, for two random variables Y_1 and Y_2 , $P(Y_1 = y_1, Y_2 = y_2) = P(Y_1 = y_2, Y_2 = y_1)$ holds for all values y_1 and y_2 . Furthermore, one of its main pillars is assigning equal mass to the probability that a post-observation falls in $n + 1$ distinct intervals created by n observations on a domain \mathbb{R} .

Assigning equal mass probability to a post-observation might seem restrictive and require further assumptions. However, when used in an imprecise probability context, it proves to be effective. The reason is that instead of relying on a single, precise probability value, uncertainty is quantified using a probability interval where optimal bounds are taken from realised observations.

In their non-parametric predictive inference (NPI) model, Augustin and Coolen [4] demonstrated the effectiveness of using lower and upper probabilities within an

$A_{(n)}$ based inference. They showed that it results in a complete framework that doesn't require additional premises. NPI is a well-established framework previously considered for system reliability problems [90], financial data modelling [9], and many more. One of its main features is the ability to fully adapt to acquired data, which makes it attractive for the problem of decision-making and learning under uncertainty.

In NPUI, observations are restricted to utility values. As opposed to NPI, instead of using a domain \mathbb{R} , these observations are bound to the interval $[0, 1]$. This is to avoid having an outcome infinitely better or worse than other alternatives. The values 0 and 1 are, respectively, the worst and best utilities of two actual or 'hypothetical' outcomes. Furthermore, the mechanism assumes that pre-observed utility values within a collection of outcomes are exchangeable. Those outcomes are considered to be sensibly grouped under a particular taxonomic category, *e.g.* sports, computer brands, etc.

Consider a set of ordered known utilities $u_{(1)}, u_{(2)} \dots u_{(n)}$, such that $0 < u_{(i)} < u_{(i+1)} < 1$. These utility values split interval $[0, 1]$ into $n + 1$ intervals. Based on assumption $A_{(n)}$, a novel outcome whose utility is exchangeable with the existing known utilities has a probability $\frac{1}{n + 1}$ of falling in one of these intervals. NPUI then extends this to lower and upper utility previsions.

Let \hat{u}_{new} denote a pre-observed utility value. Its lower and upper previsions are defined as follows:

$$\underline{P}(\hat{u}_{new}) = \frac{1}{n + 1} \sum_{i=1}^n u_i, \quad (3.2)$$

$$\overline{P}(\hat{u}_{new}) = \frac{1}{n + 1} \left(1 + \sum_{i=1}^n u_i \right) = \frac{1}{n + 1} + \underline{P}(\hat{u}_{new}). \quad (3.3)$$

Equations (3.2) and (3.3) show that the difference between the lower and upper previsions of \hat{u}_{new} is $\frac{1}{n+1}$. This difference indicates that when the number of experienced outcomes increases, NPUI reduces the range of possible values that the expected utility of a novel outcome can take. Furthermore, Equations (3.2) and (3.3) suggest that in the extreme case where there are no previously experienced outcomes with known and exchangeable utilities, this expected utility can take any value in the range $(0, 1)$. In the opposite case, where the number of experienced outcomes is infinite, lower and upper previsions coincide, which indicates that the expected utility is identified.

If the observed utility of \hat{u}_{new} is u_{new} , adding it to the set of known utilities makes it fall in one of the existing intervals. Therefore, the probability of another novel outcome having its pre-observed utility, \hat{u}_{new}^* , falling in one of the updated intervals is $\frac{1}{n+2}$. Formally, this leads to the following equations.

$$\underline{P}(\hat{u}_{new}^*|u_{new}) = \frac{1}{n+2} \left(\sum_{i=1}^n u_i + u_{new} \right) = \frac{n+1}{n+2} \underline{P}(\hat{u}_{new}) + \frac{u_{new}}{n+2}, \quad (3.4)$$

$$\overline{P}(\hat{u}_{new}^*|u_{new}) = \frac{1}{n+2} \left(1 + \sum_{i=1}^n u_i + u_{new} \right) = \frac{n+1}{n+2} \overline{P}(\hat{u}_{new}) + \frac{u_{new}}{n+2}. \quad (3.5)$$

Equations (3.4) and (3.5) highlight how a novel outcome's lower and upper utility previsions update when this outcome is experienced. This updating proves to be useful in a repeated decision-making situation that includes several unfamiliar outcomes. It allows the decision-maker to improve their utility's lower and upper previsions towards exchangeable outcomes every time one is experienced.

Suggested Model Under Ignorance

This chapter presents the suggested solution to the problem of ignorance within game theory. Section 4.1 motivates the necessity of having a model that can handle cases of ignorance. Section 4.2 introduces an enhanced GRR matrix used as a proxy to reflect a risk-neutral player's lower and upper expectations towards the payoff of an unknown outcome and subsequently compute a convex set of rational strategies. Section 4.3 discusses the rationality requirements for these lower and upper expectations, whereas Section 4.4 suggests an elicitation model to estimate them. Section 4.5 proposes a non-parametric predictive utility inference framework as an updating mechanism that allows players to adjust their lower and upper expectations dynamically after observing the payoffs of unknown outcomes. Section 4.6 extends the suggested model to handle the more generic case of risk-averse players. Sections 4.7 and 4.8 introduce an approach to assess the expected payoff of a game that includes novel outcomes and the net payoff of each. Finally, in Section 4.9, we give an example by applying our model to a variant of the game 'battle of the sexes'.

4.1 Motivation

As stated by Luce and Raïffa [47], one of the main assumptions in game theory is expecting the player to know the numerical utilities of all possible outcomes involved in the decision-making process. However, this is not always possible. In practice, it is natural to consider that a player has endogenous uncertainty over an outcome's utility caused by their ignorance and not external factors, *e.g.* other players. Often, this is ignored and either an assumed utility model from a continuous domain is used, or utility values are set using introspection.

In many cases, relying on assumed utilities is not optimal. For instance, in the *temporal sure preference* principle, Chiara et al. [19] show that prior preferences represent a minimum coherence requirement to link beliefs at different time points. It suggests that if *'you have a sure preference for A over B at (future) time t. Then you should not have a strict preference for B over A now'*; that is, prior belief should be coherent with posterior preferences. Therefore, an assumed utility limits the assessment of conditional belief toward future preferences.

Instead of assuming a known utility function, the uncertain outcome should be assessed and a player's prior unknown preferences estimated. Afterwards, these estimates can be dynamically updated when the outcome is experienced. However, such a model requires an approach to represent the source of uncertainty. Walley [86] provides a compelling and straightforward methodology to do that. He models uncertainty as a gamble bound to the possibility space of an experiment. The true state of this experiment determines the gamble's uncertain reward.

The possibility space is required to have mutually exclusive states detailed enough to describe the subject's beliefs towards the domain of interest. Hence,

a *pragmatic* possibility space can be used to include theoretical and observable states. As defined by Walley [86], it is not required to be exhaustive and should consist of sufficiently important and practically possible states, *i.e.* states with a non-zero probability of happening. Since beliefs are often incomplete under uncertainty, the possibility space can evolve and get reformulated to include new pragmatic possibilities.

Models, such as multiple prior [31] or Choquet expected utility [75], were developed to solve the problem of decision-making under uncertainty, mainly within non-strategic settings. However, as per Gajdos et al. [29], these models have certain limitations. For instance, they assume extreme pessimism, *e.g.* applying the maximin criterion to the initial set of information [31]. Furthermore, they do not allow the decision-maker to incorporate acquired information or expert knowledge in their decision making-process, hence causing an inconsistency between their decisions and actual beliefs. In practice, representing decision-maker beliefs is essential, especially since Ellsberg [27] proved that decision-makers prefer better information settings. Flexible approaches have been developed to overcome these limitations. For example, as opposed to using the maximin criterion and choosing the most pessimistic solution, Troffaes [81] discusses notions of optimality, such as E-admissibility and maximality, that allow incorporating a DM's beliefs over incomplete information. Moreover, Gajdos et al. [29] developed an axiomatic approach that allows incorporating prior information, enabling the DM to maximise their minimum expected utility with respect to this information. Nevertheless, these approaches were never formally extended to cover game theory, *i.e.* strategic settings.

Bayesian theories have emerged to account for cases where a player is in a

state of incomplete information about other players' utilities over outcomes in the game. Nevertheless, these theories don't consider a player's incomplete information about their own utilities over these outcomes. This case raises different challenges and complex sequences of reciprocal expectations [33]. For instance, in Nau's [55] operational method to achieve joint coherence and common knowledge of subjective parameters in non-cooperative games that include incomplete and observable information, a dual characterisation of joint rationality is introduced by generalising Harsanyi's [33] *Bayesian* equilibrium concept. However, in a *Bayesian* equilibrium, each player is assumed to know their actual type, considered as a summary of their actions and payoffs, and is only uncertain about other players' types [33, p.1811]. This uncertainty towards the game's structure is modelled through a common prior distribution over these players' types. Nau [55] stretches this assumption further by defining belief-revealing monetary payoff functions, *i.e.* gambles, over a set of outcomes composed of players' joint strategies and the states that represent exogenous uncertainties caused by other players or nature and not by players themselves.

Although Nau's operational method is limited to exogenous uncertainties, it includes some limitations that should be considered when modelling any type of uncertainty. First, it shows that an increased number of uncertain states results in a complex model. That is primarily because a game with incomplete information is converted into several, each representing a possible state whose lower probability is set through a *belief* gamble. The complexity could increase further if these states require updating due to acquired information, *e.g.* , in repeated games. Second, it doesn't allow simultaneously including lower and upper bounds of a possible state and computing a solution set given this range; only a scalar value within

these bounds is permitted. Finally, it doesn't incorporate a mechanism for players to update or learn about an uncertain domain through experience or acquired information.

Non-probabilistic approaches have been developed to treat the problem of players' endogenous uncertainty over outcomes. For example, Chakeri et al. [17] and Astanin and Zhukovskaja [2] apply *fuzzy logic* to game theory. Ignorance over outcomes is dealt with using the notion of fuzzy games. Fuzzy logic is a computational approach described by Elkan [26] as a more generalised propositional logic that measures degrees of truth on a scale of zero to one, rather than two truth values, 'true or false'. Fuzzy logic provides a powerful toolbox to manage situations of uncertainty in many areas, such as expert systems [92], signal processing [49], image processing [32], etc.

In game theory, there are different methods to apply *fuzzy logic*. For instance, Chakeri et al. [17] use a technique based on prioritising payoffs and measuring the preference of one against the other. First, the game is defined in fuzzy parameters by translating the player's preference degrees to a *fuzzy preference relation* matrix. This allows extending the 'greater than or equal' logic to one that is more sensitive to the difference between preferences. Afterwards, a computational algorithm, *least deviation*, is run on the fuzzy preference relation matrix to return a priority weight for each alternative.

Although computational models prove to be successful in certain situations, there is still no actual probabilistic model that could help with statistical inference and the quantification of endogenous uncertainty under strategic settings. For example, how confident can a player be that an outcome will generate a particular payoff? The interest here is using probability theory to construct a scalable

model that allows computing a game’s rational solutions based on prior lower and upper previsions of uncertain domains and updating them through experience to posterior previsions. This model would enable meaningful approximations if computation is not feasible and permit the dynamic incorporation of acquired information. Furthermore, it would provide a way to express expert knowledge, which is challenging to reflect in computational models [18].

Moreover, interest also lies in examining the consistency of game theory’s axiomatic rules under endogenous uncertainty. Von Neumann and Morgenstern’s [85] normative foundations of this theory are not extended to handle this situation. The proposed axioms of behaviour were always studied with the assumption that players know the payoff of each outcome in the game.

This study also addresses the problem of poor VOI under ignorance and provides an approach to improve it. As discussed in Section 2.2, this problem is mainly caused by limited information and can lead to non-optimal decisions.

4.2 Enhanced Revealed-Rules Matrix

Consider a game where players are risk-neutral. In practice, the payoff vectors r^k ($k = 1, \dots, K$) are not always fully known. We here develop a method for constructing a state of common knowledge of the key parameters of the payoff functions. In particular, whenever the payoff $r^k(\cdot)$ for an outcome $\phi \in \Phi$ is unknown, we introduce an arbitrary reward function $f(\cdot)$, referred to as a gamble, such that $r^k(\phi) = f(\cdot)$. This approach allows modelling ignorance using a source of uncertainty representing the pragmatic possibility space a player can face when landing an outcome with an unknown payoff. If the strategic setting involves more

than one source of uncertainty, a new gamble is required for each.

Table 4.1: Modified version of ‘battle of the sexes’ with risk-neutral players - Payoff matrix.

	L	R
T	$f, 1$	$0, 0$
B	$0, 0$	$1, 2$

Example 4.2.1. *Cont’d.* In ‘battle of the sexes’, assume that the payoff of the outcome TL is unknown to player one; that is, $r^1(TL)$ is unknown. We hence utilize a gamble f to determine $r^1(TL)$: $r^1(TL) = f(\cdot)$. Table 4.1 represents the resulting payoff matrix of this modified version of ‘battle of the sexes’.

By leveraging on the work done by Nau [56], the game can be transformed into a GRR matrix. However, the existing theory doesn’t support cases of uncertainty. Therefore, the model should be extended.

Recall, in the event where player k chooses the alternative a_i^k over any other alternative a_j^k , they are practically making a bet that is equivalent to buying the payoff vector r_i^k and selling r_j^k . However, with the presence of ignorance, if the payoff $r_i^k(\phi) \in r_i^k$ or $r_j^k(\phi) \in r_j^k$ is unknown, it is replaced with a gamble.

Example 4.2.2. *Cont’d.* In ‘battle of the sexes’, replacing the unknown payoff with gamble f results in Table 4.2.

At this stage, the GRR matrix is incomplete. It requires assessing and valuing the underlying gambles. The value of a gamble is considered to be the supremum or infimum price it is bought or sold for.

Let $\underline{P}^k(f)$ and $\overline{P}^k(f)$ be respectively the lower and upper previsions chosen by player k for gamble f . This means a player would be willing to pay $\alpha(\underline{P}^k(f) - \epsilon)$

Table 4.2: Modified version of ‘battle of the sexes’ with risk-neutral players - GRR matrix with uncertain payoffs modelled as gambles.

	TL	TR	BL	BR
$r_T^1 - r_B^1$	f	-1	0	0
$r_B^1 - r_T^1$	0	0	$-f$	1
$r_L^2 - r_R^2$	1	0	-2	0
$r_R^2 - r_L^2$	0	-1	0	2

or get paid $\alpha(\overline{P}^k(f) + \epsilon)$, in exchange for an uncertain reward αf , where $\epsilon \geq 0$ and α is a small positive number. That said, the GRR matrix can be enhanced as follows:

PROPOSITION 1. *In a non-cooperative game where endogenous uncertainty over one or several outcomes exists, an enhanced form of the revealed-rules matrix is achieved when the payoff $r_i^k(\phi)$ of each outcome ϕ in the bought vector r_i^k is replaced with its lower prevision $\underline{P}^k(r_i^k(\phi))$, and the payoff $r_j^k(\phi)$ of each outcome ϕ in the sold vector r_j^k is replaced with its upper prevision $\overline{P}^k(r_j^k(\phi))$. Furthermore, the following properties apply:*

- $\forall \phi \in \Phi$, if $r_i^k(\phi)$ does not represent a gamble, $\underline{P}^k(r_i^k(\phi)) = \overline{P}^k(r_i^k(\phi)) = r_i^k(\phi)$;
- $\forall \phi \in \Phi$, if $r_j^k(\phi)$ does not represent a gamble, $\underline{P}^k(r_j^k(\phi)) = \overline{P}^k(r_j^k(\phi)) = r_j^k(\phi)$;
- if $r_i^k(\phi)$ is a sold gamble, i.e. $r_i^k(\phi) = -f$, then $\underline{P}^k(r_i^k(\phi))$ is equal to $-\overline{P}^k(-r_i^k(\phi))$;
- if $r_j^k(\phi)$ is a sold gamble, i.e. $r_j^k(\phi) = -f$, then $\overline{P}^k(r_j^k(\phi))$ is equal to $-\underline{P}^k(-r_j^k(\phi))$;

where f denotes an arbitrary payoff function defined on outcome ϕ 's possibility space.

Now, the resulting enhanced GRR matrix permits cases of ignorance.

Example 4.2.3. *Cont'd.* Table 4.3 is the resulting enhanced GRR matrix of 'battle of the sexes'.

Table 4.3: Modified version of 'battle of the sexes' - Enhanced GRR matrix, where gamble f is replaced by its lower and upper previsions.

	TL	TR	BL	BR
$r_T^1 - r_B^1$	$\underline{P}^1(f)$	-1	0	0
$r_B^1 - r_T^1$	0	0	$-\overline{P}^1(f)$	1
$r_L^2 - r_R^2$	1	0	-2	0
$r_R^2 - r_L^2$	0	-1	0	2

The enhanced GRR matrix is interpreted as a system of inequalities that returns all correlated equilibria of a game given the specified lower and upper previsions of the underlying payoffs. This essentially means that adopting a correlated strategy depends on the player's choice of value amongst the range of possible valuations an outcome can have; that is, all values between and including the gamble's lower and upper previsions. Therefore, a choice rule is required.

Houlding and Coolen [38] proposed two choice rules that rely on the decision-maker's level of pessimism. The first is based on the attitude of Extreme Pessimism and requires choosing the outcome or sequential decision path whose lower prevision is greatest. The second is based on the attitude of Extreme Optimism and requires choosing the outcome or sequential decision path whose upper prevision is greatest.

4.3 Rationality Requirements on Lower and Upper Previsions

Since buying a gamble for a price μ is equivalent to selling it for a price $-\mu$, the supremum buying price $\underline{P}^k(f)$ should equal $-\overline{P}^k(-f)$ [86]. This allows shifting from lower to upper prevision and vice versa. Therefore, it suffices to develop the model in terms of lower previsions and use upper previsions whenever deemed necessary.

As discussed earlier, lower previsions should avoid sure loss. Any lower prevision assessment that generates loss no matter a gamble's payoff should be rejected. This requirement is ascertained using Equation (4.1), which guarantees at least one outcome x of a gamble's domain \mathcal{X} generates a payoff greater or equal to zero.

$$\sup_{x \in \mathcal{X}} \sum_{f_i \in F} [f_i(x) - \underline{P}^k(f_i)] \geq 0, \quad (4.1)$$

where F is a set of gambles.

For each gamble, a player is required to establish a lower prevision that avoids a sure loss. Then, all chosen lower previsions can be assessed for coherence using the following equation:

$$\sup_{x \in \mathcal{X}} \left\{ \sum_{f_i \in F} [f_i(x) - \underline{P}^k(f_i)] - l_0 [f_0(x) - \underline{P}^k(f_0)] \right\} \geq 0, \quad (4.2)$$

where l_0 is any positive integer and f_0 is a gamble assessed for coherence.

Equation (4.2) represents the general definition of coherence, suggested by [86]. It shows that if a lower prevision is considered acceptable, any lower prevision that

generates the same or better return should also be acceptable. Furthermore, it guarantees that assessments cannot be exploited to generate sure gains; a situation known as arbitrage that can arise if a gamble f_0 has a similar return to already accepted gambles f_i , however, with a smaller lower prevision.

Given that lower previsions, \underline{P}^k , on gambles in F avoid sure loss, coherent lower previsions can be implied for any new gamble. It is a concept known as *natural extension* that is implicitly defined in Equation (4.2). Let $\underline{\mathcal{E}}^k$ denote the natural extension of \underline{P}^k on F such that for any gamble f , $\underline{\mathcal{E}}^k(f)$ is its supremum buying price implied from $\underline{P}^k(f_i)$, through linear operations. Furthermore, $\underline{\mathcal{E}}^k$ dominates \underline{P}^k on F , which allows correcting any incoherent assessment. For instance, any previously set $\underline{P}^k(f)$ that is strictly lower than the implied one is deemed incoherent. Formally, natural extension is defined as follows:

$$\underline{\mathcal{E}}^k(f) = \sup \left\{ \omega : f(x) - \omega \geq \sum_{f_i \in F} \lambda_i [f_i(x) - \underline{P}^k(f_i)] \right\}, \quad (4.3)$$

for some $\omega \in \mathbb{R}$ and $\lambda_i \geq 0$.

4.4 First Assessment and Refinement

Let \mathcal{X} denote the set of outcomes of an experiment. Consider $f(\cdot)$ defined on \mathcal{X} . Hence, the payoff $r^k(\phi)$ of an uncertain outcome becomes: $r^k(\phi) = f(x)$, where $x \in \mathcal{X}$ is random and the probability distribution over x is unknown. Therefore, there is a lack of knowledge of the payoff.

It is essential to note that \mathcal{X} is assumed free from any uncertainty. Each state $x \in \mathcal{X}$ becomes a potential consequence if the game results in an uncertain out-

come, *i.e.* a gamble. Furthermore, a player must have complete and consistent individual preferences across the domains of available gambles and known outcomes in the game. For example, if a state can arise through gamble f , gamble g or a known outcome, its payoff is expected to be the same in all cases. Although beyond the scope of this research, Jansen et al. [39] suggest two user-friendly and robust *preference systems* that could assist a player with assessing their preferences over possible states and outcomes in the game, especially when indecisive. Their method relies on a few ranking questions that allow setting ordinal preferences. Then two different approaches can be used to determine the cardinality of these preferences. The first is a *time elicitation* approach based on the player's consideration time for ranking two states. The second is a *label elicitation* approach that relies on the player to assign pre-defined labels of preference strength.

After listing rationality requirements in Section 4.3, here we discuss how to translate beliefs into initial coherent lower previsions. Under extreme ignorance, when imprecision is at its maximum, *vacuous previsions* can be used to value gambles. They are defined as $\underline{P}^k(f) = \inf_{x \in \mathcal{X}} f(x)$ and $\overline{P}^k(f) = \sup_{x \in \mathcal{X}} f(x)$ and proven by Walley [86] to be coherent, as they respect the coherence requirements listed in Section 3.1. However, using them to model prior beliefs will lead to vacuous posterior previsions.

Usually, a player would have some prior information about a gamble, which can be used to increase the accuracy of their previsions. Therefore, amongst several methods provided by the imprecise probabilities toolbox, Walley's [86] *general elicitation* can be used to improve a vacuous assessment. It allows modelling pre-existing beliefs by translating them into explicit judgements. It is by no means a complete method that could cover all practical examples. Nevertheless, it is

sufficient enough to build our model.

First, a player starts by making qualitative judgements on elementary events in \mathcal{X} . These judgements can be comparative, *e.g.* an event is more probable than the other, or classificatory, *e.g.* an event is probable. Afterwards, judgements are modelled as almost-desirable gambles. For instance, stating that outcome x_1 is probable means that a player is willing to accept x_1 with odds better than even money. This is equivalent to accepting an almost-desirable gamble $(\delta_{x_1} - \mu)$ with a price $\mu \leq \frac{1}{2}$, where δ_{x_1} is an indicator function. Let D^k denote the set of almost-desirable gambles resulting from judgements made by player k . The following is a list of judgement examples and their relevant almost-desirable gambles:

- If outcome x_1 is probable, then gamble $\delta_{x_1} - \frac{1}{2} \in D^k$;
- If outcome x_1 is λ times as probable as outcome x_2 , then gamble $\delta_{x_1} - \lambda\delta_{x_2} \in D^k$, where $\lambda \in \mathbb{R}$;
- If outcomes x_1 and x_2 are equally likely, then gamble $\delta_{x_1} - \delta_{x_2} \in D^k$ and $\delta_{x_2} - \delta_{x_1} \in D^k$.

Using this elicitation process, a player should be able to construct any judgement representing genuine belief and model it as an almost-desirable gamble denominated in units of probability currency; that is, the payoff of this gamble is the probability of an elementary event in \mathcal{X} [86, p.59].

Once the set of almost-desirable gambles D^k is established, the second stage is to use Equation (4.4) to check that it avoids a sure loss. This equation takes a

more straightforward form than Equation (4.1).

$$\sup_{x \in \mathcal{X}} \sum_{d_q \in D^k} d_q(x) \geq 0, \quad (4.4)$$

where $d_q \in D^k$ is an almost-desirable gamble. If D^k is proven to avoid sure loss, the final stage is to compute the relevant lower prevision of each gamble $f \in F$.

Let $K(D^k)$ denote a closed convex set of linear previsions of all gambles $d_q \in D^k$. It is the intersection of all closed convex half spaces determined by the mass function $(P^k(x_1), P^k(x_2), \dots, P^k(x_{|\mathcal{X}|}))$ of each $P^k(d_q)$. Since gambles in D^k are almost-desirable, $P^k(d_q) \geq 0, \forall d_q \in D^k$. The geometry of $K(D^k)$ is a polytope on the probability simplex, satisfying the set of linear constraints applied to the possibility space \mathcal{X} . It is a credal set characterised by having a finite number of extreme points.

Using the *lower envelope theorem*, a relationship between linear previsions and coherent lower previsions can be established. The theorem suggests that coherent lower previsions \underline{P}^k of gambles in D^k are none other than the lower envelope of $P^k \in K(D^k)$, which is formally reflected in the following equation:

$$\underline{P}^k(d_q) = \min\{P^k(d_q) : P^k \in K(D^k)\}. \quad (4.5)$$

Let \mathcal{E}^k denote the natural extension of D^k . It is the set of all almost-desirable gambles implied by almost-desirable gambles in D^k ; defined as

$$\mathcal{E}^k = \sum_{d_q \in D^k} \lambda_q d_q,$$

where $\lambda_q \geq 0$. \mathcal{E}^k contains all gambles f whose $P^k(f) \geq 0$.

Based on an elementary property of polyhedral cones [30, Theorem 2.13], $K = K(D^k) = K(\mathcal{E}^k)$ is the convex hull of a finite set of linear previsions, and \mathcal{E}^k contains all gambles f whose P^k belongs to this set. Furthermore, the elements of the latter can be considered the extreme points of K . Therefore, it is denoted by $ext(K)$. Now, lower previsions of gambles in F can be computed by simply taking the lower envelope of $ext(K)$. This results in Equation (4.6).

$$\underline{P}^k(f) = \min\{P^k(f) : \forall P^k \in ext(K)\}. \quad (4.6)$$

4.4.1 Alternative models

Although this work only covers the *general elicitation* model, it is possible to use other methods to assess the gamble's initial lower and upper previsions.

Ristic et al. [74] review some of the prevalent practical methods used for quantitative modelling under ignorance, specifically when observations and the available knowledge base are uncertain. Their example-rich work compares four approaches, including imprecise probabilities, and highlights their performance through testing scenarios. A detailed discussion of each approach is beyond the scope of this research. However, we will briefly highlight them to describe alternative ways, *i.e.* other than 'imprecise probabilities', used to value a gamble when information about its possibility space is scarce.

The first approach is a modified Bayesian probabilistic model. Generally, a Bayesian model quantifies available observations using a probability function and then provides inferences using Bayes; it is used to model the known unknowns. In

its simplest form, such a model might not be suited when observable information is imprecise or unavailable [74]. Nevertheless, in the literature and practice, attempts were made to modify it and overcome this limitation [25]. For example, Mahler [48, Ch.4-8] provides a compelling Bayesian approach to modelling imprecise and random information. The application of this model can be examined in Ristic's work [73].

The second approach is a possibilistic model. It is introduced by Zadeh [91] as an extension of his work on fuzzy logic (see Section 4.1). The *possibility* theory assumes that any statement not deemed impossible cannot be excluded, and this principle is known as *minimal specificity*. In the possibilistic model, a possibility function is used to represent the level of imprecision of each event in \mathcal{X} . Based on this function, a pair of lower and upper probabilities is induced.

Finally, the third approach is based on using a belief function. A belief function was initially introduced by Shafer [76] as a more practical and flexible method for modelling ignorance over knowledge, judgements, and opinions regarding a particular event. Inspired by Shafer [76] and Dempster [24], Smets et al. [77, 78, 79] developed a belief function theory, also known as the *transferable belief model* framework. This framework provides a generic toolbox that allows the classification and quantification of any type of uncertainty.

As mentioned by [74], there is no agreement on which approach is better. Nevertheless, the imprecise probabilities framework is sufficiently rich to be considered a unified uncertainty theory. However, its numerical optimisation requires significant computation.

It is important to note that, irrespective of the approach used, the resulting model should be able to represent the available information accurately [51]. There-

fore, the imprecision caused by ignorance over elementary events in \mathcal{X} should not compromise the integrity of this model.

Furthermore, the true value corresponding to an event must exist in the set of imprecise information. For example, consider an urn that might contain black or white balls, and at least two of these balls are known to be white. Any statement similar to the following represents imprecise information about the number of white and black balls in the urn: 2 to 8 balls are white; 0 to 6 balls are black; etc. Here, stating that the urn does not contain white balls would compromise the model's integrity. Hence, it leads to erroneous modelling. In practice, a modelling error, also known as a *model-mismatch* case, cannot always be avoided, especially in particular situations where available information is limited.

4.5 Dynamic Updating

In a repeated game, whenever a player reaches an outcome that has an uncertain payoff, they get to experience it. In the suggested model, this is reflected by receiving the reward of a gamble. Each time that outcome is experienced, the player develops a preference for it. This preference evolution should be reflected through an update to the lower and upper previsions of the outcome's utility. Therefore, the lower and upper previsions of the underlying gamble should be updated. Eventually, when information about the gamble's domain is complete, these previsions converge to the gamble's linear prevision, *i.e.* its fair value (refer to Section 3.1.1). This linear prevision represents the other side of the spectrum where precision is maximal, and the supremum buying price and infimum selling price of the gamble coincide such that $P(f) := \underline{P}(f) = \overline{P}(f)$ [5, p.35].

Let n be the total number of observed values of a gamble f , where f represents the uncertain payoff of an outcome ϕ . The lower and upper previsions of $r^k(\phi)$ are $\underline{P}^k(f)$ and $\overline{P}^k(f)$. When $n = 0$, these previsions are estimated using *general elicitation* or any of the methods mentioned in Section 4.4.1.

During each game-play, whenever ϕ is experienced, *i.e.* $n > 0$, a new payoff $r^k(\phi)$ is observed. This payoff is an arbitrary utility value u_n that allows the player to learn more about the uncertain outcome's expected payoff. Given such information, the NPUI updating mechanism offers a simple yet robust (refer to Section 3.3) way to improve the lower and upper previsions of $r^k(\phi)$. Nonetheless, applying it in a strategic setting requires some modifications.

First, the NPUI model is initially developed on a unit interval $[0, 1]$, where 0 and 1 are, respectively, a DM's utilities for hypothetical worst and best outcomes. Since a utility function is unique up to a positive linear transformation, here, this model is applied to the finite interval $[w, b]$, where w and b are respectively the worst and best payoffs of gamble f .

Second, in a strategic setting, the game's outcomes do not necessarily belong to the same collection or can be grouped under the same taxonomic category. Therefore, we won't adopt an exchangeability assumption across their utility values. For instance, in Example 2.4.1, outcomes 'Hockey' (sports) and 'Cinema' (art) belong to two different categories and cannot be grouped under the same taxonomy. Hence, their utility values are considered unexchangeable. It can be argued that observed outcomes whose utility values are exchangeable with gamble f might exist outside the game. In that case, such information should be reflected in the gamble's possibility space or the elicitation model discussed in Section 4.4. For example, by using the utilities of previously experienced sports

games, a player can make a judgement about the probability of having a good experience in a novel sports game. Subsequently, this judgement can be converted into an almost-desirable gamble and used in an elicitation model. As a result, the utility's elicited lower and upper previsions would reflect the player's past experiences and pre-existing beliefs. Therefore, integrating these elicited previsions into the updating mechanism allows capturing all information known to, or believed by, the player before experiencing the novel outcome. To accomplish this, we enhance Equations (3.2) and (3.3) to include an elicited component, in this case, the one devised in Equation (4.6). Equations (4.7) and (4.8) are the result of this enhancement, which is particularly useful when no observations related to gamble f are available.

PROPOSITION 2. *Let $\{u_1, \dots, u_i, \dots, u_n\}$ denote a set of known utilities, where n is the total number of observations. In a non-cooperative game, let f denote an arbitrary utility function defined on the possibility space of an uncertain outcome Φ whose utility is exchangeable with the existing known utilities. Let $\underline{P}^k(f|u_0)$ and $\overline{P}^k(f|u_0)$ denote, respectively, player k 's initial lower and upper previsions of f when no observations exist, i.e. $n = 0$. Based on assumption $A_{(n)}$, the lower and upper previsions of the pre-observed value of f are as follows:*

$$\underline{P}^k(f) = \frac{1}{n+1} \left(\underline{P}^k(f|u_0) + \sum_{i=1}^n u_i \right), \quad (4.7)$$

$$\overline{P}^k(f) = \frac{1}{n+1} \left(\overline{P}^k(f|u_0) + \sum_{i=1}^n u_i \right). \quad (4.8)$$

The elicited components in Proposition 2 are a positive linear transformation of the unit interval $[0, 1]$ of Equations (3.2) and (3.3), where 0 and 1 are, respectively,

the worst and best utilities available. Therefore, the proof provided by [38] is still applicable. Equations (4.7) and (4.8) show that when $n = 0$, *i.e.* no observations exist, the elicited previsions $\underline{P}^k(f|u_0)$ and $\overline{P}^k(f|u_0)$ are the respective lower and upper previsions. However, once ϕ is experienced, and an actual payoff is observed, uncertainty regarding its underlying gamble is partially eliminated. In this case, the lower and upper previsions of the subsequent plays are governed by the NPUI framework.

Third, the NPUI model assumes no future outcome is better (worse) or equal to the hypothetical best (worst) outcome. In our approach, a weaker assumption is used. A future outcome can have the same utility as this best (worst) outcome, *i.e.* $w \leq u_i \leq b$.

Finally, we consider that an uncertain outcome in the game should be experienced several times before formulating a proper preference towards it. Hence, the outcome ϕ might not necessarily have the same utility every time it is observed. In practice, such flexibility is required in a strategic setting, especially towards unknown outcomes. In many circumstances, a sole experience does not reflect actual preference. The player should be allowed to try an unfamiliar outcome several times and be surprised about its payoff. This can be achieved as follows:

PROPOSITION 3. *Let $\{u_1, \dots, u_i, \dots, u_n\}$ denote a set of known utilities, where n is the total number of observations. In a non-cooperative game, let f denote an arbitrary utility function defined on the possibility space of an uncertain outcome Φ whose utility is exchangeable with the existing known utilities. Based on assumption $A_{(n)}$, when a new exchangeable utility u_{n+1} is observed, player k 's lower and upper previsions of the pre-observed value of f , respectively, $\underline{P}^k(f|u_n)$ and $\overline{P}^k(f|u_n)$ are*

updated as follows:

$$\begin{aligned}\underline{P}^k(f|u_{n+1}) &= \frac{1}{n+2} \left(\underline{P}^k(f|u_0) + \sum_{i=1}^n u_i + u_{n+1} \right) \\ &= \frac{n+1}{n+2} \underline{P}^k(f|u_n) + \frac{u_{n+1}}{n+2},\end{aligned}\tag{4.9}$$

$$\begin{aligned}\overline{P}^k(f|u_{n+1}) &= \frac{1}{n+2} \left(\overline{P}^k(f|u_0) + \sum_{i=1}^n u_i + u_{n+1} \right) \\ &= \frac{n+1}{n+2} \overline{P}^k(f|u_n) + \frac{u_{n+1}}{n+2}.\end{aligned}\tag{4.10}$$

Proposition 3 leverages Equations (3.4) and (3.5) by replacing NPUI's lower and upper previsions with the previsions provided by Proposition 2, which include an elicited component. Therefore, the proof provided by [38] is still applicable. Equations (4.9) and (4.10) show that the updated lower and upper previsions are a weighted sum of their respective values before and after observing the payoff u_{n+1} . Such updating seems intuitive. If u_{n+1} falls below the assessed lower prevision, it decreases both lower and upper previsions. If it falls above the upper prevision, it increases both lower and upper previsions. However, if it falls in between, it leads to an increase in the lower prevision and a decrease in the upper prevision.

It should be noted that the weights used in Equations (4.9) and (4.10) significantly impact how new observations are handled. The increase of existing observations will have a diminishing effect on new ones, which is sensible in a repeated game context. In practice, the early experience of an uncertain outcome greatly influences future game plays.

Updated previsions should comply with rationality requirements discussed in Section 4.3. According to Augustin and Coolen [4, Theorem 1], lower and upper

previsions based on applying assumption $A_{(n)}$ to observed data are totally monotone, and this total-monotonicity leads to coherence. Hence, Equations (4.7) and (4.8) result in coherent lower and upper previsions, assuming that the elicited previsions are also coherent. Furthermore, Augustin and Coolen [4, Theorem 7] show a strong internal consistency property in the non-parametric updating mechanism, therefore, allowing the coherence argument to be extended to cover Equations (4.9) and (4.10).

4.6 Risk Aversion

As discussed in Section 3.2.2, Nau [56] also worked on a GRR matrix that considers the non-linearity of a risk-averse player's utility. Hence, it can account for the player's state-dependent marginal utility for money. However, it fails to do so when the payoffs are unknown.

Similar to the approach used in Section 4.2, we enhance this matrix to allow for ignorance over outcomes. Consider a game where players are risk-averse. As opposed to the risk-neutral case, the utility of an outcome $\phi \in \Phi$ cannot be considered its payoff $r^k(\phi)$. Instead, it's a non-linear utility function of $r^k(\phi)$, *i.e.* $U^k(r^k(\phi))$. Therefore, whenever $r^k(\phi)$ is unknown and replaced by a gamble $f(\cdot)$ ($r^k(\phi) = f(\cdot)$), the utility of ϕ is $U^k(f(\cdot))$. For example, Table 4.4 represents the payoff matrix of the modified version of 'battle of the sexes' when played by risk-averse players.

Table 4.4: Modified version of ‘battle of the sexes’ with risk-averse players - Payoff matrix.

	L	R
T	$U^1(f), U^2(1)$	$U^1(0), U^2(0)$
B	$U^1(0), U^2(0)$	$U^1(1), U^2(2)$

To transform this game into a GRR matrix, here, in the event player k chooses the alternative a_i^k over any other alternative a_j^k , they are making a bet that is equivalent to buying the utility vector U_i^k and selling U_j^k . When the payoffs $r_i^k(\phi)$ or $r_j^k(\phi)$ are uncertain, $U^k(r_i^k(\phi)) \in U_i^k$ or $U^k(r_j^k(\phi)) \in U_j^k$ are also uncertain. In this case $r_i^k(\phi)$ or $r_j^k(\phi)$ are replaced with relevant gambles.

To assess the supremum and infimum utility previsions of each gamble, let $\underline{P}^k(U^k(f))$ and $\overline{P}^k(U^k(f))$ be respectively the lower and upper previsions chosen by player k for a gamble f . Hence, player k is willing to lose $\alpha(\underline{P}^k(U^k(f)) - \epsilon)$ or gain $\alpha(\overline{P}^k(U^k(f)) + \epsilon)$ in units of utility, in exchange for an uncertain reward αf , where $\epsilon \geq 0$ and α is a small positive number. That said, a generic GRR matrix that supports ignorance and different risk attitudes can be achieved as follows:

PROPOSITION 4. *In a non-cooperative game where endogenous uncertainty over one or several outcomes exists, an enhanced, generic form of the revealed-rules matrix that takes into account players’ marginal utilities for money is achieved when the utility $U^k(r_i^k(\phi))$ of each outcome ϕ in the bought vector U_i^k is replaced with its lower prevision $\underline{P}^k(U^k(r_i^k(\phi)))$, the utility $U^k(r_j^k(\phi))$ of each outcome ϕ in the sold vector U_j^k is replaced with its upper prevision $\overline{P}^k(U^k(r_j^k(\phi)))$, and the marginal utility $\dot{U}^k(r_i^k(\phi))$ required to convert transactions to monetary units is*

replaced by $\underline{P}^k(\dot{U}^k(r_i^k(\phi)))$. Furthermore, the following properties apply:

- $\forall \phi \in \Phi$, if payoff $r_i^k(\phi)$ does not represent a gamble, $\underline{P}^k(U^k(r_i^k(\phi))) = \overline{P}^k(U^k(r_i^k(\phi))) = U^k(r_i^k(\phi))$. Furthermore, the marginal utility of this payoff is $\underline{P}^k(\dot{U}^k(r_i^k(\phi))) = \overline{P}^k(\dot{U}^k(r_i^k(\phi))) = \dot{U}^k(r_i^k(\phi))$;
- $\forall \phi \in \Phi$, if payoff $r_j^k(\phi)$ does not represent a gamble, $\underline{P}^k(U^k(r_j^k(\phi))) = \overline{P}^k(U^k(r_j^k(\phi))) = U^k(r_j^k(\phi))$;
- if $r_i^k(\phi)$ is a sold gamble, i.e. $r_i^k(\phi) = -f$, then $\underline{P}^k(U^k(r_i^k(\phi)))$ is equal to $-\overline{P}^k(U^k(-r_i^k(\phi)))$ and $\underline{P}^k(\dot{U}^k(r_i^k(\phi)))$ is equal to $-\overline{P}^k(\dot{U}^k(-r_i^k(\phi)))$;
- if $r_j^k(\phi)$ is a sold gamble, i.e. $r_j^k(\phi) = -f$, then $\overline{P}^k(U^k(r_j^k(\phi)))$ is equal to $-\underline{P}^k(U^k(-r_j^k(\phi)))$.

Example 4.6.1. *Cont'd.* Assume the modified version of ‘battle of the sexes’ is played by risk-averse players, this results in the enhanced GRR matrix shown in Table 4.5.

Table 4.5: Modified version of ‘battle of the sexes’ with risk-averse players - Enhanced GRR matrix, where gamble f ’s utility is replaced by its lower and upper previsions.

	TL	TR	BL	BR
$\frac{U_T^1 - U_B^1}{\dot{U}_T^1}$	$\frac{\underline{P}^1(U^1(f)) - U^1(0)}{\underline{P}^1(\dot{U}^1(f))}$	$\frac{U^1(0) - U^1(1)}{\dot{U}^1(0)}$	0	0
$\frac{U_B^1 - U_T^1}{\dot{U}_B^1}$	0	0	$\frac{U^1(0) - \overline{P}^1(U^1(f))}{\dot{U}^1(0)}$	$\frac{U^1(1) - U^1(0)}{\dot{U}^1(1)}$
$\frac{U_L^2 - U_R^2}{\dot{U}_L^2}$	$\frac{U^2(1) - U^2(0)}{\dot{U}^2(1)}$	0	$\frac{U^2(0) - U^2(2)}{\dot{U}^2(0)}$	0
$\frac{U_R^2 - U_L^2}{\dot{U}_R^2}$	0	$\frac{U^2(0) - U^2(1)}{\dot{U}^2(0)}$	0	$\frac{U^2(2) - U^2(0)}{\dot{U}^2(2)}$

4.6.1 Generic Rationality Requirements on Lower and Upper Previsions

In Sections 4.2 and 4.3, utilities are generated based on a gamble's possible outcomes with the assumption that players have a state-independent linear utility for the amount of commodity gained or lost from each outcome. For example, if an outcome returns double the payoff, it is considered twice better or worse. This assumption is not valid for the risk-averse case and should be relaxed. Utilities are not necessarily linear to the payoffs. Therefore, we consider a more generic model to assess rationality requirements in this section.

The majority of statistical work focuses on inference as opposed to decision-making. This inference requires a probability assessment without the need for any utility assessment. Similarly, the imprecise probabilities framework separates *probability* from *utility*. Utility theory is only used to give a behavioural interpretation of probability [86, p.25]. As a result, statistical reasoning and inference can be conducted irrespective of a decision maker's attitude towards risk. Therefore, as Walley [86, p.10] states, there are no restrictions on the domains of definition of lower and upper previsions. An imprecise probability model can be defined on arbitrary spaces and extended to larger domains using *natural extension*.

In Section 4.3, the consistency requirements on lower and upper previsions are defined for an arbitrary domain [86, p.53]. In the risk-neutral case, this domain is a linear function f that represents the payoff of a gamble. However, in the more generic case of risk aversion, it is a function that describes the utility of that payoff, which is not necessarily linear. Therefore, the domain becomes $U^k(f)$, where $f \in F$. As a consequence, the lower and upper previsions, respectively,

$\underline{P}^k(U^k(f))$ and $\overline{P}^k(U^k(f))$ avoid sure loss when they don't cause the player to lose utility no matter the outcome of the gamble. Any lower prevision assessment that fails to comply with this requirement should be rejected. Formally, this requirement is ascertained using Equation (4.11), which guarantees at least one outcome to generate a utility that is greater or equal to zero.

$$\sup_{x \in \mathcal{X}} \sum_{f_i \in F} [U^k(f_i(x)) - \underline{P}^k(U^k(f_i))] \geq 0, \quad (4.11)$$

where F is a set of gambles.

A player is required to establish lower previsions that avoid a sure utility loss. As opposed to Equation (4.1), here, we are referring to a player's 'utility' previsions for gambles and not payoff previsions. When all these lower previsions are tested for loss aversion, they can be assessed for coherence using the following equation:

$$\sup_{x \in \mathcal{X}} \left\{ \sum_{f_i \in F} [U^k(f_i(x)) - \underline{P}^k(U^k(f_i))] - l_0 [U^k(f_0(x)) - \underline{P}^k(U^k(f_0))] \right\} \geq 0, \quad (4.12)$$

where l_0 is a positive integer and f_0 is a gamble assessed for coherence.

Equation (4.12) shows that if a lower prevision is considered acceptable, any lower prevision that generates the same or better utility should also be acceptable. It also guarantees that assessments cannot be exploited to generate sure utility gains; that is, riskless opportunities where a player can increase their utility. Therefore, if a gamble f_0 has a similar utility payoff to already accepted gambles f_i , however, with a smaller lower prevision, it should be rejected.

As seen in Section 4.3, coherent lower monetary-payoff previsions can be im-

plied through natural extension from a given set of lower previsions that avoid sure loss. The same applies to ‘utility’ assessments. However, in this case, $\underline{\mathcal{E}}$ extends \underline{P}^k on $U^k(F)$ such that for any gamble f , $\underline{\mathcal{E}}^k(U^k(f))$ is its supremum buying price implied from $\underline{P}^k(U^k(f_i))$ through linear operations. Therefore, for the risk-averse case, a natural extension is defined as follows:

$$\underline{\mathcal{E}}^k(U^k(f)) = \sup \left\{ \omega : U^k(f(x)) - \omega \geq \sum_{f_i \in F} \lambda_i [U^k(f_i(x)) - \underline{P}^k(U^k(f_i))] \right\}, \quad (4.13)$$

for some $\omega \in \mathbb{R}$, and $\lambda_i \geq 0$.

4.7 Game Expectation Under Uncertainty

In a game that includes ignorance, some payoffs are unknown. Therefore, using the payoff matrix to compute a correlated strategy’s expected utility is impossible, and Equation (2.5) is no longer valid. In this case, the enhanced GRR matrix could be used to compute this expected utility.

The enhanced GRR matrix returns the expected utility based on a player’s lower and upper utility previsions of each unknown outcome. As seen in Section 3.2.2, in the event player k chooses alternative a_i^k over any other alternative a_j^k , they are making a transaction equivalent to $U_i^k - U_j^k$. The expected utility of this choice is the dot product of this transaction’s resulting vector and the correlated strategy $\vec{\rho}$. Therefore, a player’s expected utility for the game is an aggregation of all expected utilities resulting from each possible choice. Formally, in a two-player game, this translates to the following.

PROPOSITION 5. *In a two-player game’s revealed-rules matrix, player k ’s ex-*

pected utility is

$$E^k(\vec{\rho}) = \sum_{a_i^k \in \Lambda} [U_i^k - U_j^k] \cdot \vec{\rho}, \forall j \neq i, \quad (4.14)$$

where U_i^k and U_j^k are respectively the bought and sold utility vectors of when player k chooses alternative a_i^k over any other alternative a_j^k , and $\vec{\rho}$ is the correlated strategy of the game.

Example 4.7.1. In the modified version of ‘battle of the sexes’, player one’s expected utility for a given correlated strategy $\vec{\rho}$ can be assessed by applying Equation (4.14) to the payoff matrix in Table 4.4. This is done as follows:

$$\begin{aligned} E^1(\vec{\rho}) = & [\underline{P}^1(U^1(f)) - U^1(0)] \times \rho_{TL} + [U^1(0) - U^1(1)] \times \rho_{TR} \\ & + [U^1(0) - \bar{P}^1(U^1(f))] \times \rho_{BL} + [U^1(1) - U^1(0)] \times \rho_{BR}. \end{aligned}$$

4.8 Outcome’s Expected Net Payoff

An outcome’s net payoff is the aggregate of all players’ payoffs for that outcome. For example, outcome ϕ ’s net payoff is $r^1(\phi) + \dots + r^K(\phi)$.

When particular players are uncertain about an outcome and hence ignorant about its payoff $r^k(\phi)$, it is challenging to compute this outcome’s net payoff before observing it. However, the net payoff’s estimated lower prevision can be calculated in this case. This imprecise probabilistic approach requires determining each player’s lower prevision of the relevant payoff.

As seen in Sections 3.2.2 and 4.6, the event e_i^k indicates that player k would trade any payoff vector U_j^k ($j \neq i$) for U_i^k . This trade-off translates to the transaction $U_i^k - U_j^k$. Therefore, for each outcome $\phi \in \Phi$, player k is making a bet whose

payoff is $U_i^k(\phi) - U_j^k(\phi)$. This payoff results from an exchange of wealth between players or the production/destruction of wealth in the game.

Moreover, transaction $U_i^k - U_j^k$ signifies that player k is buying the vector U_i^k , which contains the alternative i , and selling U_j^k . Hence, player k is only exposed to U_i^k , and their final payoff must be one of its elements.

Given that an outcome ϕ is the result of each player's chosen alternative, it is the intersection of all players' bought vectors. Therefore, a more generic approach to computing an outcome's net payoff is achieved through players' bought vectors. If an element in these vectors is uncertain, it can be replaced by its lower prevision $\underline{P}^k(U^k(r_i^k(\phi)))$, or $-\overline{P}^k(-U^k(r_i^k(\phi)))$ if $r_i^k(\phi)$ is a sold gamble. Formally:

PROPOSITION 6. *In a non-cooperative game, the net payoff of any outcome ϕ is*

$$\underline{P}^1(U^1(r^1(\phi))) + \dots + \underline{P}^K(U^K(r^K(\phi))), \quad (4.15)$$

where $\underline{P}^k(U^k(r^k(\phi)))$ is player k 's lower utility prevision. If an outcome's payoff $r^k(\phi)$ is not a gamble, $\underline{P}^k(U^k(r^k(\phi))) = \overline{P}^k(U^k(r^k(\phi))) = U^k(r^k(\phi))$.

Example 4.8.1. *Cont'd. In the modified version of 'battle of the sexes', Table 4.6 represents the bought payoff vector of each alternative a risk-neutral player can choose. Using these vectors and Equation 4.15, the net payoff of each outcome in this game is listed in Table 4.7.*

Table 4.6: Modified version of ‘battle of the sexes’ - Bought payoff vectors.

	TL	TR	BL	BR
r_T^1	$\underline{P}^1(f)$	0	0	0
r_B^1	0	0	0	1
r_L^2	1	0	0	0
r_R^2	0	0	0	2

Table 4.7: Modified version of ‘battle of the sexes’ - Outcomes net payoffs.

TL	$\underline{P}^1(f) + 1$
TR	0
BL	0
BR	3

4.9 Example

In the modified version of the game ‘battle of the sexes’ discussed in Section 4.2, consider that alternatives T and L stand for going to a hockey game, whereas B and R stand for going to the cinema. Furthermore, assume that player one is not familiar with hockey. Hence, their preference towards it is uncertain and replaced by gamble f .

The pragmatic possibility space \mathcal{X} can be defined as $\mathcal{X} = \{\text{Good}(G), \text{Neutral}(N), \text{Bad}(B)\}$. It represents the states a player could experience by going to the hockey game. Therefore, it can include any practically possible state. The gamble f is assigned the following payoffs $f = \{G:2, N:1, B:0\}$.

Based on the player's attitudes toward risk, the payoff matrices in Tables 4.1 and 4.4 are transformed into enhanced GRR matrices. For the risk-neutral case, the relevant GRR matrix is reflected in Table 4.3, where $\underline{P}^1(f)$ and $\overline{P}^1(f)$ are, respectively, player one's lower and upper previsions of f . Whereas for the risk-averse case, it is reflected in Table 4.5, where $\underline{P}^1(U^1(f))$ and $\overline{P}^1(U^1(f))$ are, respectively, player one's lower and upper previsions of $U^1(f)$. These matrices are still incomplete and require an assessment of the uncertain component's lower and upper previsions.

4.9.1 First Assessment

To assess gamble f 's lower and upper previsions, player one can rely on previous experiences, which might not be related to hockey, to make qualitative judgements on elementary events in \mathcal{X} . Afterwards, these judgements can be converted to almost-desirable gambles under the set D^1 . For example:

- Since they generally like sports, having a good experience is probable. This corresponds to a gamble $d_1 = \delta_G - \frac{1}{2} \in D^1$;
- Since they rarely had a bad experience at sports games in the past, having a bad experience is improbable. This corresponds to a gamble $d_2 = \frac{1}{2} - \delta_B \in D^1$;
- Since they usually like sports more than cinema, a good experience is at least as probable as a neutral one. This corresponds to a gamble $d_3 = \delta_G - \delta_N \in D^1$;
- The odds against a neutral experience are no more than 3 to 1. This corre-

sponds to a gamble $d_4 = \delta_N - \frac{1}{3} \in D^1$.

Applying Equation (4.4) to the set of gambles $D^1 = \{d_1, d_2, d_3, d_4\}$ shows that it avoids a sure loss. Since the linear prevision of each almost-desirable gamble in D^1 is greater or equal to zero and is determined by its mass function $(P^1(G), P^1(B), P^1(N))$, a credal set $K(D^1)$ can be built. It is the intersection of the following half-spaces.

$$K(D^1) = \left\{ \begin{array}{l} P^1(d_1) = P^1(G) - \frac{1}{2} \geq 0 \\ P^1(d_2) = \frac{1}{2} - P^1(B) \geq 0 \\ P^1(d_3) = P^1(G) - P^1(N) \geq 0 \\ P^1(d_4) = P^1(N) - \frac{1}{3} \geq 0 \\ P^1(G), P^1(B), P^1(N) \geq 0 \\ P^1(G) + P^1(B) + P^1(N) = 1 \end{array} \right.$$

The probability simplex in Figure 4-1 shows $K(D^1)$ and its corresponding linear previsions on the possibility space \mathcal{X} . The equilateral triangle has a height of one, and the probability of each state is identified with perpendicular distances from each side of it. The hyperplane of each gamble in D^1 cuts the simplex into a half-space. The resulting area, coloured in red, is a polyhedron that represents $K(D^1)$. Its intersections are the extreme points, $\text{ext}(K(D^1)) = \{(\frac{2}{3}, 0, \frac{1}{3}), (\frac{1}{2}, \frac{1}{6}, \frac{1}{3}), (\frac{1}{2}, 0, \frac{1}{2})\}$. The coherent lower and upper probabilities $(\underline{P}^1, \overline{P}^1)$ of each state in \mathcal{X} are the lower and upper envelopes of $\text{ext}(K(D^1))$. Hence, $(\frac{1}{2}, \frac{2}{3})$ for ‘good’, $(0, \frac{1}{6})$ for ‘bad’, and $(\frac{1}{3}, \frac{1}{2})$ for ‘neutral’. The coherent lower and upper previsions of f are, respectively, the minimum and maximum previsions of all extreme points in

$\text{ext}(K(D^1))$.




Figure 4-1: Example's credal set.

4.9.2 Risk Neutral Use Case

Using Equation (4.6), the coherent lower prevision of gamble f is computed as follows:

$$\begin{aligned}
 \underline{P}^1(f) &= \min\{P^1(f) : \forall P^1 \in \text{ext}(K(D^1))\} \\
 &= \min\left\{ \left(\frac{2}{3} \times f(G) + 0 \times f(B) + \frac{1}{3} \times f(N) \right), \right. \\
 &\quad \left(\frac{1}{2} \times f(G) + \frac{1}{6} \times f(B) + \frac{1}{3} \times f(N) \right), \\
 &\quad \left. \left(\frac{1}{2} \times f(G) + 0 \times f(B) + \frac{1}{2} \times f(N) \right) \right\} \\
 &= \min\{1.66, 1.33, 1.5\} = 1.33 .
 \end{aligned}$$

Similarly, the coherent upper prevision of gamble f is computed as follows:

$$\begin{aligned}
\bar{P}^1(f) &= \max\{P^1(f) : \forall P^1 \in \text{ext}(K(D^1))\} \\
&= \max\left\{ \left(\frac{2}{3} \times f(G) + 0 \times f(B) + \frac{1}{3} \times f(N) \right), \right. \\
&\quad \left(\frac{1}{2} \times f(G) + \frac{1}{6} \times f(B) + \frac{1}{3} \times f(N) \right), \\
&\quad \left. \left(\frac{1}{2} \times f(G) + 0 \times f(B) + \frac{1}{2} \times f(N) \right) \right\} \\
&= \max\{1.66, 1.33, 1.5\} = 1.66 .
\end{aligned}$$

Table 4.8: Modified version of ‘battle of the sexes’ with risk-neutral players - Resulting enhanced GRR matrix.

	TL	TR	BL	BR
$r_T^1 - r_B^1$	1.33	-1	0	0
$r_B^1 - r_T^1$	0	0	-1.66	1
$r_L^2 - r_R^2$	1	0	-2	0
$r_R^2 - r_L^2$	0	-1	0	2

Replacing lower and upper previsions in Table 4.3 with their relevant values, returns the enhanced GRR matrix in Table 4.8. Figure 4-2 shows that the correlated equilibria polytope of this matrix is a heptahedron with six vertices, listed in Table 4.9. Vertices two and six are pure Nash equilibria that sit at the intersection between the polytope and the simplex representing all probability distributions on outcomes, *i.e.* the tetrahedron. The remaining vertices are correlated equilibria.

It should be noted that the mixed Nash equilibrium of the original version of the game does not satisfy the correlated equilibria constraints of the modified one. Hence, on the inefficient frontier, the polytope does not intersect with the saddle that represents all joint probability distributions that are independent between

players. That's because the supremum buying price and infimum selling price of f are different. This price mismatch shows that the GRR matrix reveals information that is not obvious by just looking at the payoff matrix. Especially under ignorance, when players might have two different buy and sell values for a specific payoff.

Optimal solutions for this game sit on the edge connecting TL and BR . Choosing one of them depends on the player's choice rule (refer to Section 4.2). For instance, an extremely pessimistic player who chooses the outcome or sequential decision path with the greatest lower prevision considers $r^1(TL) = \underline{P}^1(f) = 1.33$ [38]. This choice results in an optimal equilibrium $\rho_{TL} = 0.7518$ and $\rho_{BR} = 0.2482$, with an expected game payoff of 1.248 for both players. However, an extremely optimistic player who chooses the outcome or sequential decision path with the greatest upper prevision considers $r^1(TL) = \overline{P}^1(f) = 1.66$. This choice results in an optimal equilibrium $\rho_{TL} = 0.601$ and $\rho_{BR} = 0.399$, with an expected payoff of 1.39 for both players.

Table 4.9: Modified version of 'battle of the sexes' with risk-neutral players - First assessment vertices.

	TL	TR	BL	BR
Vertex 1	0.429	0	0.215	0.356
Vertex 2	1	0	0	0
Vertex 3	0.294	0.392	0.118	0.196
Vertex 4	0.334	0.444	0	0.222
Vertex 5	0.273	0.363	0.137	0.227
Vertex 6	0	0	0	1





Figure 4-2: Modified version of 'battle of the sexes' with risk-neutral players - First assessment polytope.

4.9.3 Risk Averse Use Case

Consider players are risk averse and have an exponential utility function of the form $U^k(r^k(\phi)) = 1 - \exp(-LN(\sqrt{2}) \times r^k(\phi))$, where $LN(\sqrt{2})$ is an example of a risk aversion parameter that reflects the player's risk tolerance. In practice, this parameter can be assessed using Arrow-Pratt's measure $c = \frac{U''(\cdot)}{U'(\cdot)}$ [1] [68], the *certainty equivalent* principle (e.g. [87]), or any other relevant method. Using the defined utility function, player one's utility for each possible payoff of gamble f is $U^1(f) = \{U^1(f(G)) : 0.5, U^1(f(N)) : 0.29, U^1(f(B)) : 0\}$.

Using Equation (4.6), the coherent lower *utility* prevision of gamble f is computed as follows:

$$\begin{aligned} \underline{P}^1(U^1(f)) &= \min\{P^1(U^1(f)) : \forall P^1 \in \text{ext}(K(D^1))\} \\ &= \min\left\{ \left(\frac{2}{3} \times U^1(f(G)) + 0 \times U^1(f(B)) + \frac{1}{3} \times U^1(f(N)) \right), \right. \\ &\quad \left(\frac{1}{2} \times U^1(f(G)) + \frac{1}{6} \times U^1(f(B)) + \frac{1}{3} \times U^1(f(N)) \right), \\ &\quad \left. \left(\frac{1}{2} \times U^1(f(G)) + 0 \times U^1(f(B)) + \frac{1}{2} \times U^1(f(N)) \right) \right\} \\ &= \min\{0.43, 0.346, 0.395\} = 0.346 . \end{aligned}$$

Similarly, the coherent upper *utility* prevision of gamble f is computed as follows:

$$\begin{aligned} \overline{P}^1(U^1(f)) &= \max\{P^1(U^1(f)) : \forall P^1 \in \text{ext}(K(D^1))\} \\ &= \max\{0.43, 0.346, 0.395\} = 0.43 . \end{aligned}$$

The local marginal utility of money, *i.e.* the first derivative of the given utility

function, is $\dot{U}^k(r^k(\phi)) = LN(\sqrt{2})exp(-LN(\sqrt{2}) \times r^k(\phi))$. Therefore, $\dot{U}^1(f) = \{\dot{U}^1(f(G)) : 0.173, \dot{U}^1(f(N)) : 0.245, \dot{U}^1(f(B)) : 0.346\}$, and the lower and upper previsions of $\dot{U}^1(f)$ are computed as follows:

$$\begin{aligned} \underline{P}^1(\dot{U}^1(f)) &= \min\{P^1(\dot{U}^1(f)) : \forall P^1 \in ext(K(D^1))\} \\ &= \min\left\{\begin{aligned} &\left(\frac{2}{3} \times \dot{U}^1(f(G)) + 0 \times \dot{U}^1(f(B)) + \frac{1}{3} \times \dot{U}^1(f(N))\right), \\ &\left(\frac{1}{2} \times \dot{U}^1(f(G)) + \frac{1}{6} \times \dot{U}^1(f(B)) + \frac{1}{3} \times \dot{U}^1(f(N))\right), \\ &\left(\frac{1}{2} \times \dot{U}^1(f(G)) + 0 \times \dot{U}^1(f(B)) + \frac{1}{2} \times \dot{U}^1(f(N))\right) \end{aligned}\right\} \\ &= \min\{0.197, 0.168, 0.209\} = 0.168, \end{aligned}$$

and

$$\begin{aligned} \overline{P}^1(\dot{U}^1(f)) &= \max\{P^1(\dot{U}^1(f)) : \forall P^1 \in ext(K(D^1))\} \\ &= \max\{0.197, 0.168, 0.209\} = 0.209. \end{aligned}$$

Table 4.10: Modified version of ‘battle of the sexes’ with risk-averse players - Resulting enhanced GRR matrix.

	TL	TR	BL	BR
$\frac{U_T^1 - U_B^1}{\dot{U}_T^1}$	2.05	-0.84	0	0
$\frac{U_B^1 - U_T^1}{\dot{U}_B^1}$	0	0	-1.24	1.19
$\frac{U_L^2 - U_R^2}{\dot{U}_L^2}$	1.19	0	-1.44	0
$\frac{U_R^2 - U_L^2}{\dot{U}_R^2}$	0	-0.84	0	2.89

Replacing lower and upper previsions in Table 4.5 with their relevant values

returns the enhanced GRR matrix in Table 4.10.

Figure 4-3 shows that the enhanced GRR matrix's correlated equilibria polytope is a heptahedron with six vertices, listed in Table 4.11. Vertices two and six are pure Nash equilibria that sit at the intersection between the polytope and the simplex representing all the probability distributions on outcomes, *i.e.* the tetrahedron. The remaining vertices are correlated equilibria. On the inefficient frontier, the polytope intersects with the saddle representing all joint probability distributions independent between players. Therefore, as proven by Nau [57, Proposition 1], the points of intersections on the surface of the polytope are mixed Nash Equilibria. That is because a Nash equilibrium satisfies at least one of the GRR matrix constraints or one of the non-negativity constraints with equality. Hence, this constraint coupled with the equality constraint *i.e.* $\vec{\rho}'1 = 1$ determines a face of the correlated equilibria polytope, which indicates that a Nash equilibrium cannot be an interior point of the latter.

Similar to the risk-neutral case, optimal solutions for this game sit on the edge connecting TL and BR . Choosing one of them depends on the player's level of pessimism. For instance, an extremely pessimistic player would consider TL 's utility as the gamble's lower utility prevision, *i.e.* $U^1(r^1(TL)) = \underline{P}^1(U^1(f)) = 0.346$. Furthermore, they would consider the marginal utility as $\dot{U}^1(r^1(TL)) = \underline{P}^1(\dot{U}^1(f)) = 0.168$. This results in an optimal equilibrium $\rho_{TL} = 0.793$ and $\rho_{BR} = 0.206$, where both players have an expected game utility of 0.335 and an expected monetary payoff of 1.2, denominated in the game's payoff currency.

Conversely, an extremely optimistic player would assume that $U^1(r^1(TL)) = \overline{P}^1(U^1(f)) = 0.43$. This results in an optimal equilibrium $\rho_{TL} = 0.6$ and $\rho_{BR} = 0.4$, where both players have an expected game utility of 0.374 and an expected

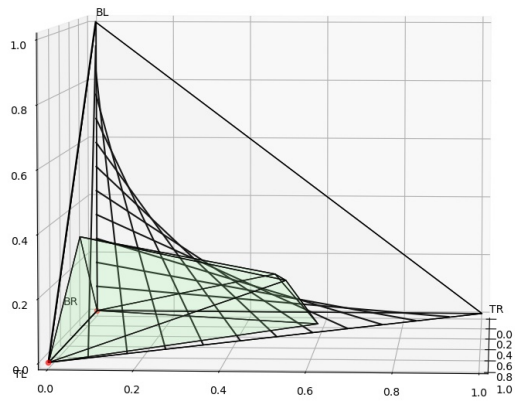



Figure 4-3: Modified version of 'battle of the sexes' with risk-averse players - First assessment polytope.

monetary payoff of 1.4, denominated in the game’s payoff currency.

Table 4.11: Modified version of ‘battle of the sexes’ with risk-averse players - First assessment vertices.

	TL	TR	BL	BR
Vertex 1	0.372	0	0.307	0.320
Vertex 2	1	0	0	0
Vertex 3	0.24	0.587	0	0.171
Vertex 4	0.206	0.504	0.141	0.147
Vertex 5	0.195	0.476	0.16	0.167
Vertex 6	0	0	0	1

Table 4.12: Dynamic updating applied to three different scenarios.

n	Scenario 1			Scenario 2			Scenario 3		
	$u_n = f(x)$	$\underline{P}^1(f)$	$\overline{P}^1(f)$	$u_n = f(x)$	$\underline{P}^1(f)$	$\overline{P}^1(f)$	$u_n = f(x)$	$\underline{P}^1(f)$	$\overline{P}^1(f)$
0	-	1.33	1.66	-	1.33	1.66	-	1.33	1.66
1	2	1.665	1.83	0	0.665	0.83	0	0.665	0.83
2	2	1.776	1.886	0	-	-	2	1.11	1.22
3	2	1.832	1.915	0	-	-	2	1.332	1.415

4.9.4 Repeated Game Use Case

Sections 4.9.2 and 4.9.3 showcase two examples of how the initial set of correlated equilibria is computed for a game that includes ignorance. These examples cover use cases of risk-neutral and risk-averse players. Furthermore, they are limited to the first play of a game or a one-shot game. In this section, we cover the repeated game use case and show how dynamic updating improves lower and upper previsions whenever a player experiences an uncertain outcome. Subsequently, these improved previsions return improved correlated equilibria.

The dynamic updating model introduced in Section 4.5 does not depend on the players’ risk attitudes and can be applied to risk-averse or risk-neutral cases.

To avoid redundancy, only the latter will be treated in this example. Table 4.12 illustrates three different scenarios of applying dynamic updating to a sequence of plays. It shows that whenever outcome TL has a new payoff, the underlying gamble's lower and upper previsions are updated. Consequently, this triggers an update to the optimal correlated strategy.

Scenario one considers that the first time a play ends with an outcome TL , player one enjoys it more than all other outcomes. Hence, the payoff of gamble f is 2. Since this observation falls above the elicited $\bar{P}^1(f)$, the updated lower and upper previsions increase in value. Furthermore, $\underline{P}^1(f|u_1) = 1.665$ indicates that an extremely optimistic or pessimistic player expects TL 's payoff to be the highest amongst all other payoffs still. If player one keeps getting a payoff of 2 for TL , the lower and upper previsions will eventually converge towards 2. In this case, the enhanced GRR matrix returns the correlated equilibria polytope in Figure 2-1.

Scenario two shows the opposite case. After experiencing outcome TL , player one considers it the worst outcome in the game. Hence, the payoff is 0. Since this observation falls below the elicited $\underline{P}^1(f)$, both updated previsions decrease in value. An upper prevision $\bar{P}^1(f|u_1) = 0.83$ is strictly smaller than the payoff player one gets from outcome BR . Therefore, an extremely optimistic or pessimistic player would stop choosing the alternative T and settles for a correlated strategy of $\rho_{TL} = 0$ and $\rho_{BR} = 1$, *i.e.* the pure Nash equilibrium. However, note that if only the upper prevision is higher than BR 's payoff, an extremely optimistic player would still consider the alternative T .

Scenario three shows the case where player one has different experiences related to outcome TL . The first time it is observed, its payoff is 0. This payoff decreases the upper prevision to $\bar{P}^1(f|u_1) = 0.83$, which is enough for a player

to discard the alternative T . However, in practice, an optimistic player can still explore the outcome TL as a trial attempt. In that case, if the second observation has a payoff of 2, $\underline{P}^1(f|u_1)$ and $\overline{P}^1(f|u_2)$ increase in value and both become strictly higher than BR 's payoff. As discussed by Houlding and Coolen [37], trial attempts are frequently observed in daily life. It is natural for a DM to experiment with novel outcomes, *e.g.* try a new item on a food menu. Therefore, having a model that adjusts to new observations resulting from a trial attempt is beneficial. Furthermore, it allows customising Houlding and Coolen's [38] choice rules (refer to Section 4.2) to include an element of exploring novel outcomes and exploiting them when relevant.

Special Cases

This chapter discusses the application of the suggested model to special cases of uncertainty. In particular, Section 5.1 covers the topic of extreme ignorance, where a player has completely no information about a specific outcome. A Monte Carlo simulation is used to demonstrate the effectiveness of the suggested model under this case. Section 5.2 covers the case of multiple sources of uncertainty in a game and highlights global rationality requirements across these sources. Finally, Sections 5.3 and 5.4 discuss the repercussion of adopting an enhanced GRR matrix and its effect on a game's structure, precisely *zero-sum* and *symmetry* properties.

5.1 Extreme Ignorance

To illustrate our model under a case of extreme ignorance, we consider a variant of the game 'Matching Pennies'. Von Neumann and Morgenstern [85] describe the classic version as a game where two players simultaneously and independently select 'Heads' or 'Tails' each and then uncover a penny. If their selections match, then player two must give a penny to player one. Otherwise, player one gives a penny to player two. However, we modify the game here so that player two gives

player one an arbitrary reward generated by a gamble g . This gamble consists of drawing a ball from an urn, and depending on its colour, the following rewards are generated: 0 for red, 1 for black, and 2 for green.

Table 5.1: Modified version of ‘matching pennies’ - Payoff matrix.

	T	H
T	$g, -g$	$-1, 1$
H	$-1, 1$	$g, -g$

We assume that information is symmetric across players throughout the game-play and that no information about the composition of the urn is available to them. Therefore, vacuous lower and upper previsions are used; $\underline{P}^1(g) = \underline{P}^2(g) = 0$ and $\overline{P}^1(g) = \overline{P}^2(g) = 2$. Under these circumstances, whether the game is played or not depends on the adopted choice rule.

Using the pessimist/optimist decision rules, if any of the players is a pessimist, they will use the lower previsions of underlying gambles to assess expected payoffs. For outcomes TT and HH , they would only expect to lose utility when playing this game. Hence, they don’t have any incentive to play it.

However, if players one and two are optimists, they would expect, respectively, a payoff of $\overline{P}^1(g) = 2$ and $\overline{P}^2(-g) = -\underline{P}^2(g) = 0$ for both of these outcomes. Therefore, when no reward is yet observed for gamble g , the game’s expectation for both optimists players is 0.5, and its mixed Nash equilibrium is $(1/2 \text{ H}, 1/2 \text{ T}) \times (1/2 \text{ H}, 1/2 \text{ T})$.

Consider that the urn contains one black, one green, and two red balls. Applying the suggested dynamic updating algorithm (refer to Section 4.5) to 1000

simulations of 200 plays each, returns an average lower prevision of $\underline{P}^1(g|u_{200}) = \underline{P}^2(g|u_{200}) = 0.74$ and an upper prevision of $\overline{P}^1(g|u_{200}) = \overline{P}^2(g|u_{200}) = 0.75$.




Figure 5-1: Modified version of ‘matching pennies’ - Average lower and upper previsions generated using 1000 simulations with 200 plays each

Figure 5-1 shows how gamble g ’s lower and upper previsions converge towards its linear prevision; that is, the actual expected payoff $P(g) = 0.75$, which is unknown to the players. The dynamic updating will influence each player’s estimate of the game’s expected payoff. As seen in Figure 5-2, on average, player one’s expectation becomes negative after the fifth play, giving them no incentive to keep playing the game. In contrast, player two’s expectation is always positive. This indicates that player two has an advantage over player one, which is expected since the urn contains two red balls.

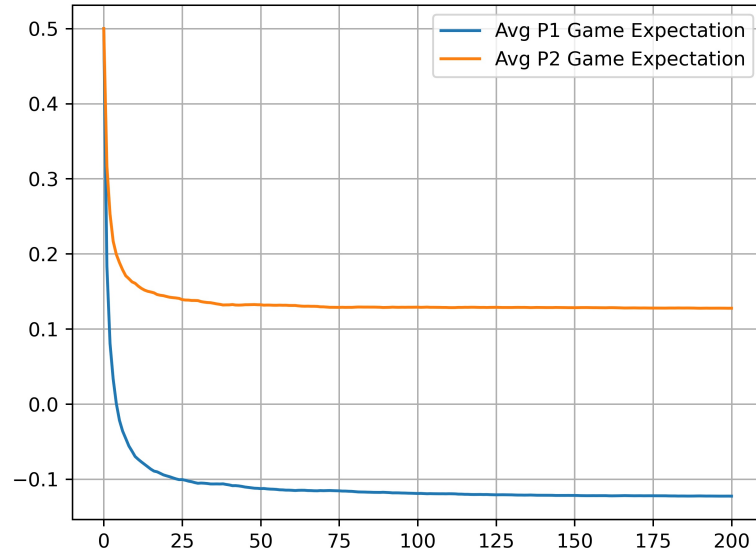


Figure 5-2: Modified version of ‘matching pennies’ - Players’ average game expectations generated using 1000 simulations, with 200 plays each.

As Section 4.1 highlights, computational approaches exist for games with incomplete information. It is noteworthy that these approaches do not necessarily rely on a gamble nor a pragmatic possibility space, *i.e.* \mathcal{X} , to model an uncertain domain. For instance, in this variant of ‘Matching Pennies’, when the game results with outcome TT or HH , both players are exposed to three states: lose, win, or draw. Therefore, outcomes TT and HH involve fuzziness caused by the uncertainty of their payoffs. In certain fuzzy solution concepts (*e.g.* [45] and [46]), this uncertainty is reflected through an interval of payoffs. Such games are known as matrix games with interval payoffs. The literature offers several methods to solve these games and compute their optimal strategies. However, some of the existing methods are limited to specific cases, *e.g.*, zero-sum games. Moreover, some methods are shown to be incorrect. Verma and Kamit [83] revisited the incorrect

methods and highlighted their flaws. For example, they showed that Li et al. 's [45] method is mathematically inaccurate and that Liu and Kao's [46] method fails to return optimal solutions. Verma and Kamit [83] proposed alternative solutions. Nevertheless, these solutions are still subject to the limitations of computational models discussed in Section 4.1.

5.2 Multiple Sources of Ignorance

So far, we have dealt with one source of uncertainty modelled as an arbitrary reward function. However, in some games, a player can be exposed to multiple sources of ignorance. For example, in Section 4.9, player one may have never been to the cinema as well; they are ignorant about this outcome. Therefore, they have to define a second gamble on a new pragmatic possibility space representing the theoretical and observable states they might get from going to the cinema.

The suggested model is scalable and supports multiple sources of uncertainty. In this case, each gamble has its possibility space. However, these gambles' lower and upper previsions should not be limited to *local assessments*, *i.e.* assigning lower and upper previsions to each gamble independently. Instead, *structural assessments* that consider the relationship between lower (upper) previsions across all gambles should be used.

Structural assessments allow a player to express judgements about global properties related to their beliefs. They cover assessments of independence and symmetry. Therefore, they should be integrated into the player's imprecise probability model to ensure that local assessments satisfy global properties.

We briefly discuss the structural judgement of independence, particularly *epis-*

temic independence. A more in-depth study of imprecise probabilities' different types of structural judgements is provided by Walley [86, ch.9]. Further discussions can also be found in [5].

Epistemic independence is used as a conditioning criterion to infer or improve lower and upper previsions assessments between gambles with different domains. Based on Kolmogorov's conditional probability approach [42], a player would consider an outcome $x \in \mathcal{X}$ *epistemically irrelevant* to outcome $y \in \mathcal{Y}$ when they assume learning the true state of x doesn't affect their beliefs towards y . Formally,

$$P(y|x) = \frac{P(x, y)}{P(x)} = P(y), \quad (5.1)$$

$\forall (x, y) \in \mathcal{X} \times \mathcal{Y}$ and $P(x) \neq 0$.

In an imprecise probabilities context, such assessment is helpful as it allows inferring or adjusting coherent lower previsions of gambles on \mathcal{X} from existing ones on \mathcal{Y} .

Epistemic irrelevance is an asymmetric concept. If x is epistemically irrelevant to y , this doesn't necessarily suggest that y is epistemically irrelevant to x . However, when x and y are epistemically irrelevant to each other, they are then *epistemically independent*.

Example 5.2.1. *Let \mathcal{X} and \mathcal{Y} represent respectively the pragmatic possibility spaces of states a player k could experience when going to a hockey game or cinema. They are defined as follows: $\mathcal{X} = \mathcal{Y} = \{Good(G), Neutral(N), Bad(B)\}$.*

Consider the state player k gets from going to hockey epistemically irrelevant to the state they get from going to the cinema. That indicates the lower (upper) previsions of gambles on \mathcal{X} have no impact on lower (upper) previsions on \mathcal{Y} .

Therefore no information on \mathcal{X} can be used to adjust local assessments on \mathcal{Y} , which for an outcome $y = G$ translates to the following equations:

$$\underline{P}^k(y = G|x = G) = \underline{P}^k(y = G|x = B) = \underline{P}^k(y = G|x = N)$$

and

$$\overline{P}^k(y = G|x = G) = \overline{P}^k(y = G|x = B) = \overline{P}^k(y = G|x = N).$$

If the state player k gets from hockey impacts the state they get from going to the cinema, then x and y are not epistemically irrelevant. If this player considers a bad experience in hockey leads to a probably good or neutral experience in the cinema, the resulting equations are as follows: $\underline{P}^k(y = G|x = B) = 0.5$, $\underline{P}^k(y = N|x = B) = 0.5$, and $\underline{P}^k(y = B|x = B) = 0$.

5.3 Ignorance within Zero-Sum Games

Games can be either *zero-sum* or *non-zero-sum*. In non-zero-sum games (*e.g.* ‘battle of the sexes’), outcomes’ net payoffs (see Section 4.8) are not zero nor necessarily constant [85, p.46]. However, in zero-sum games, these net payoffs are always zero. Hence, one player’s loss is another player’s win, and assets (*e.g.* money) are not produced externally or destroyed during the game.

Zero-sum games represent extreme states of competition between players [44]. They help model particular strategic settings, mainly in Economics and political sciences. Furthermore, they are used by Von Neumann and Morgenstern [85] as a medium for constructing a theory for all games.

Like non-zero-sum games, zero-sum games could include cases of ignorance

where a player is uncertain about an outcome's payoff. In these situations, the enhanced GRR matrix could represent this game as *non-zero-sum*. That is because a player's elicited lower and upper previsions don't necessarily match other players' assessments for an uncertain outcome. Hence, players do not agree on the expected amount that should be paid or received for this outcome.

For example, consider the variant of 'matching pennies' described in Section 5.1. In this game, when the selections of both players match, they exchange an arbitrary reward generated by gamble g . Therefore, both players have the exact source of uncertainty. However, each can assess gamble g 's lower and upper previsions differently. Hence players are not expecting to pay or receive the same payoff for the relevant outcomes. As a consequence, the zero-sum property of the game is not respected.

Table 5.2: Modified version of 'matching pennies' with risk-neutral players - Enhanced GRR matrix.

	TT	TH	HT	HH
$r_T^1 - r_H^1$	$\underline{P}^1(g) + 1$	$-1 - \overline{P}^1(g)$	0	0
$r_H^1 - r_T^1$	0	0	$-1 - \overline{P}^1(g)$	$\underline{P}^1(g) + 1$
$r_T^2 - r_H^2$	$-\overline{P}^2(g) - 1$	0	$1 + \underline{P}^2(g)$	0
$r_H^2 - r_T^2$	0	$1 + \underline{P}^2(g)$	0	$-\overline{P}^2(g) - 1$

Table 5.3: Modified version of 'matching pennies' with risk-neutral players - Finalised enhanced GRR matrix.

	TT	TH	HT	HH
$r_T^1 - r_H^1$	1	-3	0	0
$r_H^1 - r_T^1$	0	0	-3	1
$r_T^2 - r_H^2$	-3	0	2	0
$r_H^2 - r_T^2$	0	2	0	-3

Example 5.3.1. Table 5.2 shows the enhanced GRR matrix of the modified ver-

sion of ‘matching pennies’. Consider player one is completely ignorant about the content of the urn used to determine gamble g ’s reward. Therefore, their initial vacuous lower and upper previsions are $\underline{P}^1(g) = 0$ and $\overline{P}^1(g) = 2$. However, player two knows that this urn doesn’t contain any red balls. Hence, the only possible rewards are 1 (black) and 2 (green). Such information would influence player two’s assessments. For instance, it can push them to set their initial vacuous lower and upper previsions to $\underline{P}^2(g) = 1$ and $\overline{P}^2(g) = 2$.

As seen in Table 5.3, players’ lower and upper previsions result in a GRR matrix that doesn’t reflect a zero-sum game. For each relevant outcome, the payoff of a player’s bet is not expected to be entirely paid or received by the other player. For instance, events e_H^1 and e_H^2 lead to outcome HH . Under this outcome, player one’s bet has a payoff of 1, whereas player two’s bet has a payoff of -3 . Hence, wealth destruction is expected, which contradicts zero-sum game properties.

When strategic settings involve uncertainty over outcomes, dropping the zero-sum game assumption is inevitable, especially when players have different sources of information to assess payoffs. Asymmetric information makes it challenging for players to agree on the expected payoff of the unknown domain.

In a zero-sum structure, the net payment of each outcome in the game is expected to be zero. However, when outcomes are uncertain, these payments are estimates. Therefore, if the estimated net payment of these outcomes is 0, the game is zero-sum. This net payment is estimated using Equation (4.15), which allows for ignorance over outcomes.

Example 5.3.2. *The modified version of ‘matching pennies’ is zero-sum if:*

$$\underline{P}^1(r^1(TT)) + \underline{P}^2(r^2(TT)) = 0$$

$$\underline{P}^1(g) + \underline{P}^2(-g) = 0$$

$$\underline{P}^1(g) - \overline{P}^2(g) = 0$$

$$\underline{P}^1(g) = \overline{P}^2(g)$$

and

$$\underline{P}^1(r^1(HH)) + \underline{P}^2(r^2(HH)) = 0$$

$$\underline{P}^1(g) + \underline{P}^2(-g) = 0$$

$$\underline{P}^1(g) - \overline{P}^2(g) = 0$$

$$\underline{P}^1(g) = \overline{P}^2(g).$$

These results indicate that players must agree on the bought and sold prices of gamble g .

Alternatively, using multiple enhanced GRR matrices for analysing the game can enforce a zero-sum structure. In this case, each matrix would assume that elicited previsions across all players align with a single player’s expectations. For example, two enhanced GRR matrices can be used in the modified version of ‘matching pennies’. The first matrix assumes that player two’s upper prevision $\overline{P}^2(g)$ matches player one’s lower prevision $\underline{P}^1(g)$. The second matrix assumes that player one’s lower prevision $\underline{P}^1(g)$ matches player two’s upper prevision $\overline{P}^2(g)$.

5.4 Ignorance within Symmetric Games

Numerous strategic settings are represented through symmetric games. As discussed by Cao et al. [16], that is primarily because symmetry often simplifies the analysis of a problem without compromising its essence. Hence why symmetric games are popular in the game theory literature and treated as benchmarks in various application areas. Furthermore, Papadimitriou and Roughgarden [62] show that symmetry in games simplifies representational and computational challenges.

Von Neumann and Morgenstern [84] consider a game *ordinary symmetric* or *totally symmetric* if all players have the same alternatives and payoff function. Cao et al. [16] extend this definition using elementary concepts of the *permutation group theory*. They first define *ordinary player symmetry* as follows:

Definition 2. Let $\Sigma(K)$ denote the set of permutations of the set of players K . A permutation $\pi \in \Sigma(K)$ is an *ordinary player symmetry* if for each outcome $\phi \in \Phi$, $\pi(\phi) \in \Phi$ and $r^k(\phi) = r^{\pi(k)}(\pi(\phi)), \forall k \in K$.

Subsequently, they define *ordinary symmetry* as follows:

Definition 3. A game is *ordinary symmetric* if all permutations $\pi \in \Sigma(K)$ are *ordinary player symmetry*.

Example 5.4.1. Consider the game ‘prisoner’s dilemma’, devised by Merrill Flood and Melvin Dresher in 1950, then named and formalised by Albert Tucker [67] [82]. In this game, two persons are arrested and suspected of a crime. With the available evidence, they are convicted for a minor offence. The police take each person to a separate interrogation room and try to acquire further evidence by making them

confess. Therefore, they have two alternatives, confess (C) or keep quiet (Q). The game's set of possible outcomes whose payoffs are denominated in utility units is as follows:

- If both players keep quiet, they are charged with a minor offence. To introduce uncertainty, the utility of this outcome is considered unknown and replaced with a gamble f .
- If only one player confesses and the other keeps quiet, the former does not incur any charges, and the latter is charged two utility units.
- If both players confess, they are charged one utility unit each.

The resulting payoff matrix of this modified version of 'prisoner's dilemma' is shown in Table 5.4.

Table 5.4: Modified version of 'prisoner's dilemma' with risk-neutral players - Payoff matrix.

	C	Q
C	-1, -1	0, -2
Q	-2, 0	f, f

Using Definitions 3 and 2, the modified version of 'prisoner's dilemma' is ordinary symmetric if:

- $r^1(CC) = r^2(CC)$
- $r^1(CQ) = r^2(QC)$
- $r^1(QC) = r^2(CQ)$

$$- r^1(QQ) = r^2(QQ)$$

Hence, for the uncertain outcome QQ , players one and two's expected utilities for gamble f should be in line. In particular, the players' lower previsions of f should match (see Section 4.8), i.e. $\underline{P}^1(f) = \underline{P}^2(f)$. It can be noticed that unlike 'prisoner's dilemma', 'battle of the sexes' is not ordinary symmetric in its standard and modified versions, e.g. $r^1(HH) \neq r^2(HH)$.

Definition 3 is not entirely satisfactory. For instance, a game is not considered symmetric if players have the same payoff function and alternatives have different names [16]. Therefore, Peleg et al. [66] and Sudhölter et al. [80] propose a less restrictive definition known as *name-irrelevant symmetric* games. It considers the names of the alternatives irrelevant and focuses solely on the players' symmetry and payoffs.

Cao et al. define *name-irrelevant symmetric games* by extending Definition 3 to use a set of permutations of alternatives available to each player, which is a notion suggested by Nash [54]. Furthermore, influenced by differences between the symmetry structures of 'battle of the sexes' and 'matching pennies', Cao et al. define a new symmetry class known as *renaming symmetric games*. It relies on renaming alternatives. *Name-irrelevant symmetric games* and *renaming symmetric games* are beyond the scope of this study. However, the analysis of ignorance within symmetric games, performed in Example 5.4.1, also applies to these classes of symmetry.

When uncertainty exists, players should agree on the utility previsions of the uncertain outcomes to get a symmetric structure. In repeated games, these previsions should be in line throughout each game-play. However, this can be chal-

lenging when asymmetric information exists. For example, players who hold extra information adjust their previsions accordingly, leading other players to disagree over these assessments. Another case that could lead to this situation is players using different elicitation models to assess their previsions.

Multiple enhanced GRR matrices can be used as a workaround to enforce symmetry and avoid differences between players over expected utilities. In this case, each matrix would assume that elicited previsions across all players align with a single player's expectations. For example, two enhanced GRR matrices can be used in the modified version of 'prisoner's dilemma'. The first matrix assumes that player two's lower prevision $\underline{P}^2(g)$ matches player one's lower prevision $\underline{P}^1(g)$. The second matrix assumes that player one's lower prevision $\underline{P}^1(g)$ matches player two's lower prevision $\underline{P}^2(g)$.

Utility Diagnostics Under Uncertainty

This chapter discusses the implication of ignorance on utility diagnostics and introduces an approach that allows assessing the impact of information patterns on strategic interactions. This approach relies on the enhanced GRR matrix to compute a piece of information's expected and actual value. We demonstrate it using the previously introduced variant of 'battle of the sexes'.

6.1 Value of Information

DeGroot [23] introduced the value of sample information derived from utility diagnostics. This concept was further studied in [11, 22, 72] and is limited to *decision theory*, which only considers the decision-making process in a non-strategic setting. Decision theory analyses the process and consequence of choosing an option amongst a set of alternatives. It aims at helping a DM make the best decision possible.

Houlding and Coolen [37] extend the traditional approach of valuing sample

information to include cases of uncertainty over preferences. Furthermore, they discuss trial aversion and introduce an approach to assess the DM's aversion toward decisions with high uncertainty.

Here we extend the value of sample information through utility diagnostics to game theory to assess the impact of information patterns on strategic interactions. We consider the case where a player is ignorant about a game's outcomes and uncertain about their preferences for these outcomes.

This case raises interesting questions regarding the impact of uncertainty on VOI. To answer these questions and quantify the VOI of a player's true preferences, we introduce an approach that complements the enhanced GRR matrix presented in Section 4.6. The proposed solution allows a player to value new information based on its impact on their expected utility of a game.

Let Π denote a random variable used to represent a currently unrevealed piece of information. Let \mathcal{I} denote a set of possible information statements and ι a statement in \mathcal{I} . In a non-strategic setting, the expected amount of information in Π is considered its fair utility value. It is the supremum amount of utility a subject would sacrifice to know Π . Furthermore, it is the difference between the supremum expected utility of a decision where Π is known and the supremum expected utility of a decision where Π is unknown.

The definition of the expected amount of information under a non-strategic setting should be extended to a strategic setting. In the latter, the player is not acting unilaterally. They are part of a group of self-interested entities, and each entity is trying to maximise its welfare by adopting a strategy for interaction [65]. Therefore, in a game with multiple players, the expected amount of information of Π differs between players and is relative to each player's adopted strategy. In

other words, player k 's expected amount of information of Π is the difference between the supremum expected utility of a strategy where Π is known and the supremum expected utility of a strategy where Π is unknown, while noting that these strategies don't have to be the same.

In the decision theory framework, Π is used to provide the DM with information about the likely outcome of an available decision. For instance, if r denotes a possible payoff and d denotes an available decision, there exists a relation between r , d , and the information statement ι such that the probability $P(r|d, \Pi = \iota) \neq P(r|d)$. However, as seen in Section 3.2, in game theory, when a decision is made, player k chooses an alternative a_i^k from the set Λ^k and this choice exposes them to a vector of possible payoffs r_k whose underlying outcomes include alternative a_i^k .

The outcome of a game is the result of all players' chosen alternatives. Therefore, for the set of outcomes that include player k 's alternative a_i^k , the relation between the player's payoff, the player's decision, and the information statement ι becomes:

$$P(r^k|a_i^k, \Pi = \iota) \neq P(r^k|a_i^k). \quad (6.1)$$

Equation (6.1) indicates that player k 's vector of payoffs r_i^k changes when an information statement ι is acquired. The same logic is extended to assess the relationship between the information provided by Π and the enhanced GRR matrix.

In the event player k chooses alternative a_i^k over any other alternative a_j^k , they are making a transaction equivalent to buying U_i^k and selling U_j^k . Hence, the impact of an information statement ι on this transaction's *final* payoff vector, *i.e.*

$U_i^k - U_j^k$, is as follows:

$$P(\{U_i^k - U_j^k\} | \Pi = \iota) \neq P(\{U_i^k - U_j^k\}). \quad (6.2)$$

Let $V^k(\iota)$ and $V^k(\Pi)$ denote, respectively, the value of information ι and the value of unknown information Π with respect to player k . Inspired by DeGroot, Raiffa, and Schlaifer's work [72, 23], we extend the expected value of sample information to game theory as follows:

PROPOSITION 7. *In a strategic setting, the expected value of sample information $V^k(\Pi)$ can be computed as follows:*

$$E^k[V^k(\Pi)] = E^k[U^k(\vec{\rho}_2)|\Pi] - E^k[U^k(\vec{\rho}_1)], \quad (6.3)$$

where $\vec{\rho}_1$ is an ex-ante Pareto efficient correlated strategy, $\vec{\rho}_2$ is an ex-post Pareto efficient correlated strategy, and $E^k[U^k(\vec{\rho}_2)|\Pi]$ is the game's expected utility given the information set Π .

Equation (6.3) assesses the value of sample information relative to a Pareto efficient correlated strategy that reflects the joint beliefs of all players. It is a correlated equilibrium where no player can be better off without making at least one player worse off [63].

Using a Pareto efficient correlated strategy to assess the value of sample information is advantageous. It is more specific than the correlated equilibria set and reflects a more practical approach toward games that include ignorance. That is because when particular outcomes are unknown, a piece of relevant information could create an opportunity for Pareto improvement and push players to

re-evaluate their beliefs towards the correlated strategy. Hence, they might negotiate a new Pareto efficient equilibrium. Renegotiating a correlated strategy upon acquiring information about a novel outcome is not applicable when all outcomes are known. In this case, payoffs used to agree on a correlated randomisation do not change. Therefore, players have no incentive to deviate from the agreed correlated strategy, irrespective of what information they get.

When $\Pi = \iota$ is observed, the distribution $P(\{U_i^k - U_j^k\} | \Pi = \iota)$ becomes the actual beliefs of player k upon which a correlated equilibrium should be assessed. In this case:

PROPOSITION 8. *In a strategic setting, the value $V^k(\iota)$ of an observed statement ι is as follows:*

$$V^k(\iota) = E^k[U^k(\vec{\rho}_p)|\iota] - E^k[U^k(\vec{\rho}^*)|\iota], \text{ with } \vec{\rho}^* = \arg \max \{E^k[U^k(\vec{\rho})]\} \quad (6.4)$$

where $\vec{\rho}^*$ is the pre-observation correlated strategy and $\vec{\rho}_p$ is the post-observation correlated strategy.

Equation (6.4) is an extension of the work done by Raiffa & Schlaifer [72] and DeGroot [23]. As opposed to the *ex ante* measure $E^k[V^k(\Pi)]$, $V^k(\iota)$ is an *ex post* diagnostic. Computing it requires the player to identify the optimal strategy of the game before the knowledge that $\Pi = \iota$. Therefore, the player should be able to determine the strategy $\vec{\rho}^*$ that they would have played before acquiring new information. The knowledge of $\vec{\rho}^*$ is essential after $\Pi = \iota$ becomes known, as it affects the value of $V^k(\iota)$. However, if Π is not known yet, the only relevant measure is $E^k[V^k(\Pi)]$, which does not depend on $\vec{\rho}^*$.

In non-strategic settings, DeGroot shows that more information would never

negatively impact. If a DM wishes, they can ignore it. Hence, $V^k(\iota) \geq 0$ for all $\iota \in \mathcal{I}$. However, Osborne [60] discusses that the same doesn't necessarily apply to strategic settings. In particular cases, more information could make a player worse off when strategies are not correlated. For instance, in [60, p.281], Osborne showcases a *Bayesian* game with two possible states. Each state represents a player's *type* (refer to Section 4.1) through a separate payoff matrix. In this game, player two gets a higher expected payoff from not knowing player one's type. When this type is revealed, player two's expected payoff decreases. Therefore, player two is better off not acquiring this information.

Houlding and Coolen [37] note that $E^k[V^k(\Pi)]$ and $V^k(\iota)$ are introduced by Raiffa & Schlaifer [72, Ch.4] as, respectively, the Expected Value of Sample Information (EVSI) and the Conditional Value of Sample Information (CVSI). We extended these two measures to non-strategic settings, allowing their usage in game theory. Specifically, in cases where a player is ignorant about the payoff of an outcome, hence, uncertain about its utility. We showed that EVSI relies on the following game-theoretic parameters:

- The player $k = 1, \dots, K$, where K is the total number of players;
- The sets of alternatives $\Lambda^k = \{a_1^k, \dots, a_{m^k}^k\}$, where m^k is the total number of alternatives available to player k ;
- The utility function U^k adopted by each player k ;
- The *ex ante* correlated equilibrium $\vec{\rho}_1$ that is agreed between players, irrespective of the information set Π ;
- The *ex post* correlated equilibrium $\vec{\rho}_2$ that is achieved by combining the

likelihood function $P(\Pi|U^k(\vec{\rho}_1))$ with $P(U^k(\vec{\rho}_1))$. The resulting posterior expected utility $P(U^k(\vec{\rho}_1)|\Pi)$ pushes the player to negotiate a new Pareto efficient correlated equilibrium, *i.e.* $\vec{\rho}_2$.

In the suggested model, the additive property of the classical EVSI persists. Therefore, before making any move, if the player receives multiple pieces of information simultaneously, their expected value is equivalent to the expected value of receiving them sequentially, *i.e.* one after the other. Formally, when a piece of information $\Pi_1 = \iota_1$ is revealed and the player negotiates a new Pareto efficient correlated equilibrium accordingly, the EVSI of Π_2 is as follows:

$$\begin{aligned}
E^k[V^k(\Pi_1, \Pi_2)] &= E^k[E^k[V^k(\Pi_2|\Pi_1)] + V^k(\Pi_1)] \\
&= E^k[V^k(\Pi_2|\Pi_1)] + E^k[V^k(\Pi_1)] \\
&= E^k[U^k(\vec{\rho}_3)|\Pi_2, \Pi_1] - E^k[U^k(\vec{\rho}_2)] + \\
&\quad E^k[U^k(\vec{\rho}_2)|\Pi_1] - E^k[U^k(\vec{\rho}_1)]
\end{aligned} \tag{6.5}$$

6.2 Example

In Section 4.9, we compute the correlated equilibria of the modified game ‘battle of the sexes’ and demonstrate how the enhanced GRR matrix is used when the utility of an outcome is unknown to the player. Furthermore, we use a general elicitation model to assess an unknown outcome’s initial lower and upper utility provisions. Here, we highlight how information acquired by the player can impact these initial assessments, affecting the computed set of correlated equilibria. In this example, we only consider information acquired before a player makes their move and show how the value of this information is assessed through utility diagnostics.

We previously evaluated the lower and upper previsions of gamble f 's payoff for a risk-neutral player. We assessed the probability of each event in $\mathcal{X}=\{\text{Good}(G), \text{Neutral}(N), \text{Bad}(B)\}$, then computed the lower and upper previsions as follows:

$$\begin{aligned}
\underline{P}^1(f) &= \min\{P^1(f) : \forall P^1 \in \text{ext}(K(D^1))\} \\
&= \min\left\{ \left(\frac{2}{3} \times f(G) + 0 \times f(B) + \frac{1}{3} \times f(N) \right), \right. \\
&\quad \left(\frac{1}{2} \times f(G) + \frac{1}{6} \times f(B) + \frac{1}{3} \times f(N) \right), \\
&\quad \left. \left(\frac{1}{2} \times f(G) + 0 \times f(B) + \frac{1}{2} \times f(N) \right) \right\} \\
&= \min\{1.66, 1.33, 1.5\} = 1.33
\end{aligned}$$

and

$$\begin{aligned}
\overline{P}^1(f) &= \max\{P^1(f) : \forall P^1 \in \text{ext}(K(D^1))\} \\
&= \max\left\{ \left(\frac{2}{3} \times f(G) + 0 \times f(B) + \frac{1}{3} \times f(N) \right), \right. \\
&\quad \left(\frac{1}{2} \times f(G) + \frac{1}{6} \times f(B) + \frac{1}{3} \times f(N) \right), \\
&\quad \left. \left(\frac{1}{2} \times f(G) + 0 \times f(B) + \frac{1}{2} \times f(N) \right) \right\} \\
&= \max\{1.66, 1.33, 1.5\} = 1.66
\end{aligned}$$

Let Π denote a random variable that informs player one of the probable outcomes of gamble f . For instance, it can represent an unrevealed piece of statistics about riots in hockey games. Let the set $\mathcal{I} = \{Riots, Peace\}$ represent the possible values that Π can take. Let the probability of a hockey game ending with a riot be

$P(\text{riot}) = \theta$, and the probability of peace $P(\text{peace}) = 1 - \theta$. Consider that a riot would certainly lead player one to have a bad experience. Therefore, $P(B|\text{riot}) = 1$ and $P(G|\text{riot}) = P(N|\text{riot}) = 0$. However, a peaceful game would still keep the player unsure about their preferences towards the hockey game. Hence, they still rely on the lower and upper utility previsions resulting from the extreme points in $\text{ext}(K(D^1))$ (refer to Section 4.9.1).

In Equation (6.3), to assess the expected value of Π , *i.e.* $E^1[V^1(\Pi)]$, player one needs to determine the expected utilities of the *ex-ante* and *ex-post* Pareto efficient correlated strategies, hence, compute $E^1[U^1(\vec{\rho}_2)|\Pi]$ and $E^1[U^1(\vec{\rho}_1)]$.

As seen in Section 4.9, before acquiring any information, an extremely pessimistic player one considers the payoff of outcome TL to be $r^1(TL) = \underline{P}^1(f) = 1.33$. This assessment results in an optimal equilibrium $\vec{\rho}_1 = \{\rho_{TL} = 0.75, \rho_{BR} = 0.25\}$, hence, an expected game payoff of $E^1[U^1(\vec{\rho}_1)] = 1.25$. In contrast, an extremely optimistic player one considers the payoff of outcome TL to be $r^1(TL) = \overline{P}^1(f) = 1.66$. This results in an optimal equilibrium $\vec{\rho}_1 = \{\rho_{TL} = 0.6, \rho_{BR} = 0.4\}$, with an expected utility of $E^1[U^1(\vec{\rho}_1)] = 1.39$.

Computing the conditional expected utility $E^1[U^1(\vec{\rho}_2)|\Pi]$ is not as straightforward as $E^1[U^1(\vec{\rho}_1)]$. It requires player one to identify the *ex-post* Pareto efficient correlated strategy $\vec{\rho}_2$; that is, the correlated strategy after acquiring Π . First, the player should assess the impact of Π on the initial lower and upper previsions $\underline{P}^1(f)$ and $\overline{P}^1(f)$, hence, compute $\underline{P}^1(f|\Pi)$ and $\overline{P}^1(f|\Pi)$. This can be done by assessing the following conditional probabilities for each extreme point in $\text{ext}(K(D^1))$:

- $P(G|\Pi) = P(\text{riot}) \times P(G|\text{riot}) + P(\text{peace}) \times P(G|\text{peace});$
- $P(B|\Pi) = P(\text{riot}) \times P(B|\text{riot}) + P(\text{peace}) \times P(B|\text{peace});$

$$- P(N|\Pi) = P(riot) \times P(N|riot) + P(peace) \times P(N|peace).$$

When these conditional probabilities are established for each extreme point in $\text{ext}(K(D^1))$, player one's conditional lower and upper previsions of gamble f 's payoff are computed as follows:

$$\begin{aligned} \underline{P}^1(f|\Pi) &= \min\{P^1(f|\Pi) : \forall P^1 \in \text{ext}(K(D^1))\} \\ &= \min\{P^1(G|\Pi) \times f(G) + P^1(B|\Pi) \times f(B) \\ &\quad + P^1(N|\Pi) \times f(N) : \forall P^1 \in \text{ext}(K)\} \\ &= \min\left\{\left(\frac{2(1-\theta)}{3} \times f(G) + \frac{(1-\theta)}{3} \times f(N)\right) \right. \\ &\quad \left(\frac{(1-\theta)}{2} \times f(G) + \frac{(1-\theta)}{3} \times f(N)\right) \\ &\quad \left(\frac{(1-\theta)}{2} \times f(G) + \frac{(1-\theta)}{2} \times f(N)\right) \\ &= \min\left\{\frac{5(1-\theta)}{3}, \frac{4(1-\theta)}{3}, \frac{3(1-\theta)}{2}\right\} = \frac{4(1-\theta)}{3}, \end{aligned}$$

and

$$\begin{aligned} \overline{P}^1(f|\Pi) &= \max\{P^1(f|\Pi) : \forall P^1 \in \text{ext}(K(D^1))\} \\ &= \max\left\{\frac{5(1-\theta)}{3}, \frac{4(1-\theta)}{3}, \frac{3(1-\theta)}{2}\right\} = \frac{5(1-\theta)}{3}. \end{aligned}$$

The game's conditional expected utility given $\underline{P}^1(f|\Pi)$ and $\overline{P}^1(f|\Pi)$ is computed using Equation(2.5). Subsequently, the expected value of Π is computed using Equation (6.3). The following demonstrates how this is applied using the pessimist and optimist decision rules.

If player one is extremely pessimistic, their conditional expected utility is as-

essed using the conditional lower prevision of f . Hence,

$$\begin{aligned} E^1[U^1(\vec{\rho}_2)|\Pi] &= \underline{P}^1(f|\Pi) \times \rho_{2TL} + r^1(TR) \times \rho_{2TR} \\ &\quad + r^1(BL) \times \rho_{2BL} + r^1(BR) \times \rho_{2BR} \\ &= \frac{4(1-\theta)}{3} \times \rho_{2TL} + \rho_{2BR}. \end{aligned}$$

Then, the EVSI of this player is:

$$\begin{aligned} E^1[V^1(\Pi)] &= E^1[U^1(\vec{\rho}_2)|\Pi] - E^1[U^1(\vec{\rho}_1)] \\ &= \frac{4(1-\theta)}{3} \times \rho_{2TL} + \rho_{2BR} - 1.25 \end{aligned}$$

Figure 6-1 shows player one's conditional lower prevision of gamble f as a function of θ , *i.e.* the probability of a riot. Furthermore, it shows the conditional expected payoff of player one given the conditional lower prevision $\underline{P}^1(f|\Pi)$. For instance, if the probability of a riot is $\theta < 0.25$, the conditional lower prevision is $\underline{P}^1(f|\Pi) > 1$, and the conditional expected payoff $E^1[U^1(\vec{\rho}_2)|\Pi] > 1$. In this case, the payoff prevision for outcome TL is higher than the payoff of BR . Hence, player one would still consider a correlated strategy that includes both outcomes.

If the probability of a riot is $\theta = 0.25$, the conditional lower prevision is $\underline{P}^1(f|\Pi) = 1$, and the conditional expected payoff is $E^1[U^1(\vec{\rho}_2)|\Pi] = 1$. In this case, the payoff prevision for outcome TL is the same as BR . Hence, player one is indifferent between TL and BR . Both of these outcomes generate the same utility. Since player two's payoff is $r^2(BR) = 2$, a Pareto efficient correlated strategy would be $\rho_{2BR} = 1$.

If the probability of a riot is $\theta > 0.25$, the conditional lower prevision is $\underline{P}^1(f|\Pi) < 1$, and the conditional expected payoff $E^1[U^1(\vec{\rho}_2)|\Pi] = 1$. In this

case, player one would always prefer the outcome BR over TL as it returns a higher payoff. Hence, the Pareto efficient correlated strategy is $\rho_{2BR} = 1$.




Figure 6-1: Player one's expected payoff and initial assessment of gamble f 's lower prevision, given the probability of a riot. The expected payoff is based on a Pareto efficient correlated strategy and on the assumption that player one is extremely pessimistic.

For the opposite case, where player one is extremely optimistic, the conditional upper prevision $\bar{P}^1(f|\Pi)$ is used instead to compute the expected conditional payoff. Hence,

$$\begin{aligned}
 E^1[U^1(\vec{\rho}_2)|\Pi] &= \bar{P}^1(f|\Pi) \times \rho_{2TL} + r^1(TR) \times \rho_{2TR} \\
 &\quad + r^1(BL) \times \rho_{2BL} + r^1(BR) \times \rho_{2BR} \\
 &= \frac{5(1-\theta)}{3} \times \rho_{2TL} + \rho_{2BR}.
 \end{aligned}$$

Then, the EVSI of this player is:

$$\begin{aligned} E^k[V^k(\Pi)] &= E^k[U^k(\vec{\rho}_2)|\Pi] - E^k[U^k(\vec{\rho}_1)] \\ &= \frac{5(1-\theta)}{3} \times \rho_{2TL} + \rho_{2BR} - 1.39. \end{aligned}$$

Figure 6-2 shows player one's conditional upper prevision f as a function of θ . Furthermore, it shows the conditional expected payoff of player one given the conditional upper prevision $\bar{P}^1(f|\Pi)$. For instance, if the probability of a riot is $\theta < 0.4$, the conditional upper prevision is $\bar{P}^1(f|\Pi) > 1$, and the conditional expected payoff $E^1[U^1(\vec{\rho}_2)|\Pi] > 1$. In this case, the payoff prevision for outcome TL is higher than the payoff of BR . Hence, player one would still consider a correlated strategy that includes both outcomes. Given player one's choice rule, this result is expected. Even though the probability of a riot is high, the player is optimistic about the outcome of TL and is still willing to go to a hockey game as long as $\theta < 0.4$.

If the probability of a riot is $\theta = 0.4$, the conditional upper prevision is $\bar{P}^1(f|\Pi) = 1$, and the conditional expected payoff is $E^1[U^1(\vec{\rho}_2)|\Pi] = 1$. In this case, the payoff prevision for outcome TL is the same as BR . Hence, player one is indifferent between TL and BR . Both of these outcomes generate the same utility. Since player two's payoff is $r^2(BR) = 2$, a Pareto efficient correlated strategy would be $\rho_{2BR} = 1$.

If the probability of a riot is $\theta > 0.4$, the conditional upper prevision is $\underline{P}^1(f|\Pi) < 1$, and the conditional expected payoff $E^1[U^1(\vec{\rho}_2)|\Pi] = 1$. In this case, player one would always prefer the outcome BR over TL as it returns a higher payoff. Hence, the Pareto efficient correlated strategy is $\rho_{2BR} = 1$.

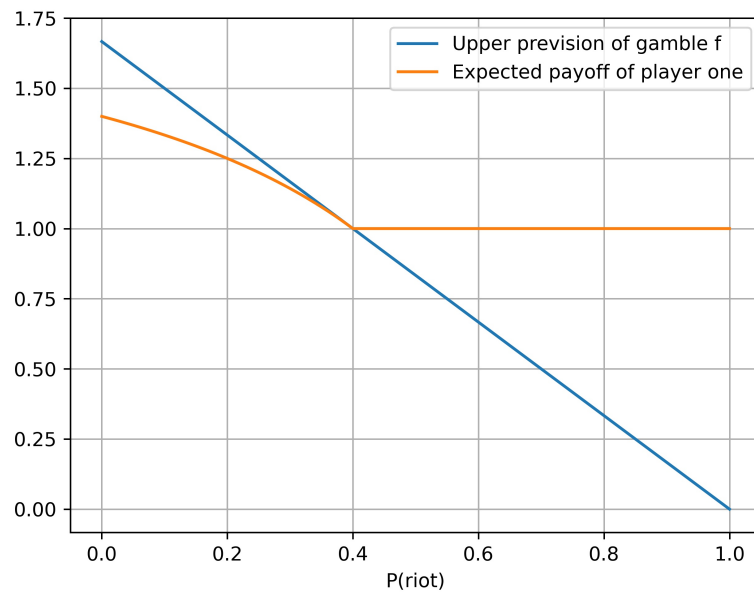


Figure 6-2: Player one’s expected payoff and initial assessment of gamble f ’s upper prevision, given the probability of a riot. The expected payoff is based on a Pareto efficient correlated strategy and on the assumption that player one is extremely pessimistic.

Figure 6-3 shows the EVSI of Π , *i.e.* $E^1[V^1(\Pi)]$, as a function of θ . It analyses the impact of the provided information on the expected value of this information. Here, we analyse the impact of the probability of riot, θ , on the EVSI of Π , $E^1[V^1(\Pi)]$. For instance, if the player is informed that hockey games are always peaceful and riots are uncommon, the initially adopted correlated strategy is not impacted. Hence, the player’s expected payoff doesn’t change. However, the initial correlated strategy is significantly impacted if the player is informed that riots are probable, *i.e.* $\theta = 0.5$. Whether the player is extremely optimistic or pessimistic, this information pushes them to negotiate a new correlated strategy. In this case, they would only consider going to the cinema. The new correlated strategy causes their expected payoff to change by -0.25 for the ‘pessimistic’ case and -0.4 for

the ‘optimistic’ case.

Furthermore, Figure 6-3 shows that any information resulting in $\theta > 0.25$ is treated the same by a pessimistic player and has the same EVSI. That is because it causes the player to switch to a new correlated strategy that excludes the outcome TL . The same analysis applies to the optimistic player when $\theta > 0.4$. Moreover, when $\theta < 0.3$, the EVSI of an extremely optimistic player is greater than the EVSI of an extremely pessimistic player. This observation is explained by the fact that an optimistic player’s conditional expected payoff of outcome TL is higher than that of a pessimistic player. However, when $\theta \geq 0.3$, the pessimistic player’s EVSI becomes higher. That is because the information that makes a player ignore the outcome TL has a more significant impact on an optimistic player than on a pessimistic player. This information reduces the optimistic player’s expected utility by -0.4 and the pessimistic player’s expected utility by -0.25 .

The value of a particular statement ι in $\mathcal{I} = \{Riots, Peace\}$ can be identified using the following steps. First, we compute the game’s expected utility given the initial correlated strategy and given the relevant statement $\iota \in \mathcal{I}$. For example, consider that probability of a riot is $\theta < 0.25$ for an extremely pessimistic player one or $\theta < 0.4$ for an extremely optimistic player. Furthermore, consider that the *ex-ante* Pareto efficient strategy results in players one and two going to the hockey match. During the hockey match, riots start and player one’s payoff is zero, *i.e.* $r^1(TL) = 0$. Figure 6-4 shows $E^1[U^1(\vec{\rho}_2)|riot]$ as a function of θ ; that is, player one’s expected payoff given the *ex-ante* Pareto efficient strategy, $\vec{\rho}_2$, and given that a riot started.

Second, we compute the game’s expected payoff given the new negotiated correlated strategy and the relevant statement $\iota \in \mathcal{I}$. For instance, the riot event

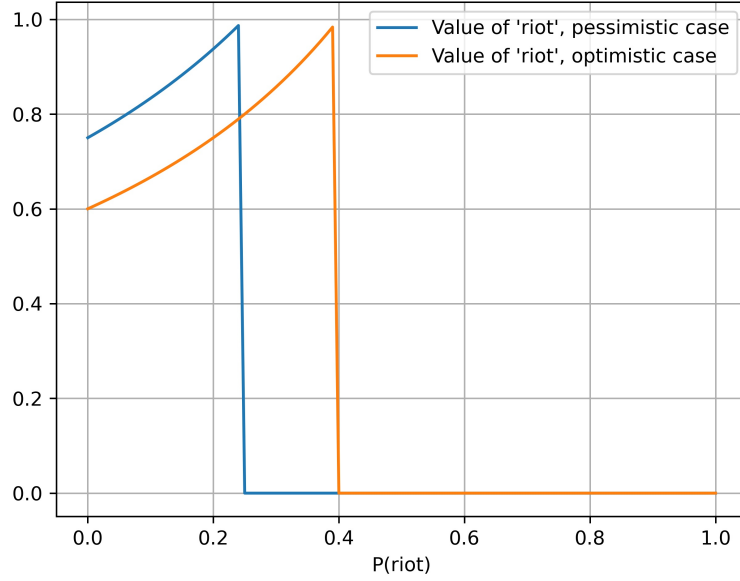


Figure 6-3: Expected value of sample information of both a pessimistic and optimistic player one, given the probability of a riot.

pushes players to negotiate a new correlated strategy $\vec{\rho}_p$. In this example, consider that they decide to stop going to hockey games. Hence, player one's expected payoff given the *ex-post* correlated strategy and given that a riot starts is $E^1[U^1(\vec{\rho}_p)|riot] = 1$.

Finally, using Equation 6.4, the value of a 'riot' is computed as follows:

$$\begin{aligned}
 V^1(riot) &= E^1[U^1(\vec{\rho}_p)|riot] - E^1[U^1(\vec{\rho}_2)|riot] \\
 &= 1 - \rho_{2BR}.
 \end{aligned}$$

Figure 6-5 shows the value of a 'riot' as a function of θ . For a pessimistic player, when the probability of a riot is $\theta = 0$, the value of a 'riot' statement is $V^1(riot) = 0.75$, whereas when $\theta = 0.24$, the value of a 'riot' statement is

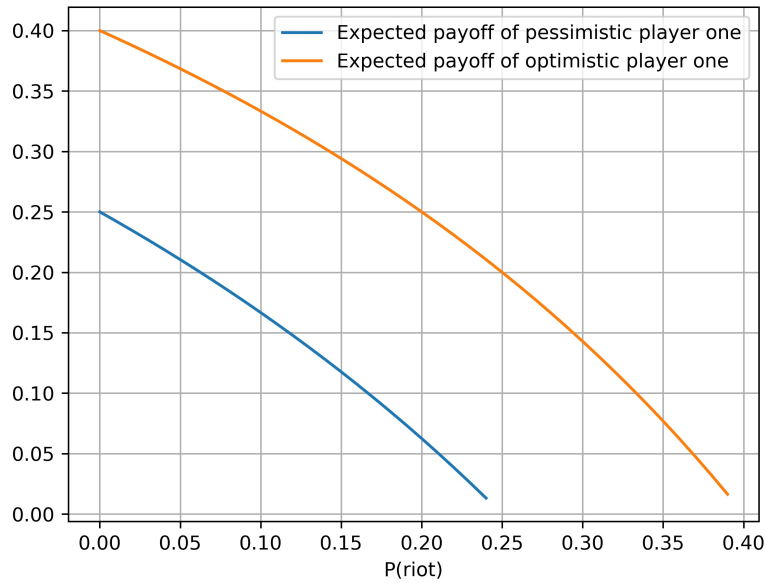


Figure 6-4: Expected payoff of a pessimistic and optimistic player one given the *ex-ante* Pareto efficient strategy and given that a riot starts.

$V^1(riot) = 0.98$. These values are expected, a statement informing the player of a high probability of riot would help them negotiate a better correlated equilibrium. Then, if a riot happens, its impact on this player would be minimal. However, a statement informing the player of a low probability of riot, *e.g.* $\theta = 0$, is less valuable. In this case, if a riot happens, the player will lose a significant amount of utility.

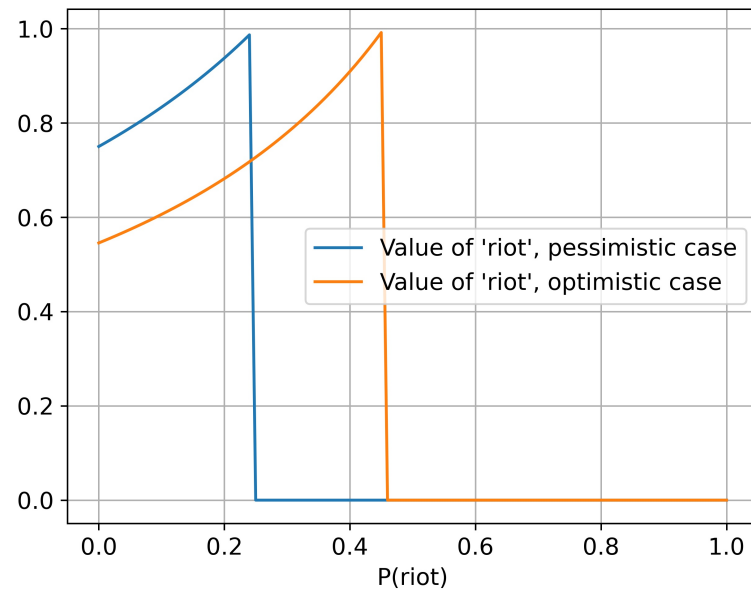


Figure 6-5: Value of a 'riot' for a pessimistic and optimistic player.

Discussion and Conclusions

7.1 Discussion

The proposed enhanced GRR matrix, dynamic updating model, and utility diagnostics approach could help expand the existing scope of application of game theory and allow it to include cases of ignorance over outcomes. Moreover, we believe this solution could be extended to cooperative games. For instance, in oligopoly pricing [3] or repeated partnerships [71], players are assumed to know the utility of each outcome and the set of correlated distributions over these outcomes. However, this strong assumption could be relaxed using an enhanced GRR matrix.

Consider Radner's [70] work on enhancing the decentralised decision-making process in an organisation. He studies repeated partnership games in which players cannot observe each other's strategies. In his example, two players contribute separate efforts to an enterprise. The combined effort of all players leads the enterprise to succeed or fail. Hence, $\mathcal{X} = \{\text{Success}(S) = 1, \text{Fail}(F)=0\}$. Players choose their effort simultaneously without being able to monitor each other's choices. The probability of success is considered $P(x = 1) = \min(e^1 + e^2, 1)$, where e^1 and e^2

are the individual efforts of players one and two. Furthermore, the utility payoff of a player k is $f^k(x, e^k) = x - \lambda(e^k)^2$, where $\lambda > 0$, and λa^k represents the *disutility of effort*.

In this game, when a player contributes high effort, they know it is more likely to yield success. However, they don't know the exact probability of a failure when all players choose to do the same. Therefore, players are uncertain about their expected payoffs. An enhanced GRR matrix can help assess this uncertainty and compute the set of correlated equilibria. Consider that a high effort(H) is when $e^k > 0.5$ and a low effort(L) is when $e^k \leq 0.5$. Let $\Lambda^1 = \Lambda^2 = \{H, L\}$ denote the alternatives available to each player. Hence, the set of possible outcomes is $\Phi = \{HH, HL, LH, LL\}$. Outcomes HH and LL represent respectively a sure success and a sure loss. However, both outcomes, HL and LH, represent either a success or a loss. Assuming $\lambda = 1$, the payoff functions of these outcomes can be modelled as follows:

- $r^k(HH) = f_1^k(e^k) = 1 - (e^k)^2$, this function represents the payoff each player k gets when they and the other player choose to contribute high effort. Since, for the outcome HH, each player has to provide a minimum effort of 0.5, the lower and upper previsions of $r^k(HH)$ are $\underline{P}^k(r^k(HH)) = 0.75$ and $\overline{P}^k(r^k(HH)) = 0$ (assuming that the maximum effort is 1);
- $r^k(HL) = f_2^k(x, e^k) = x - (e^k)^2$ and $r^k(LH) = f_3^k(x, e^k) = x - (e^k)^2$, these functions represent the payoff each player k gets when they choose to contribute high effort while the other player contributes low effort, and vice versa. The lower and upper previsions of $r^k(HL)$ are $\underline{P}^k(r^k(HL)) = \underline{P}^k(x) - (\underline{P}^k(e^k))^2$ and $\overline{P}^k(r^k(HL)) = \overline{P}^k(x) - (\overline{P}^k(e^k))^2$. The same logic applies to

payoff $r^k(LH)$. Hence, $\underline{P}^1(r^1(HL)) = \underline{P}^2(r^2(LH)) = 0.25$, $\overline{P}^1(r^1(HL)) = \overline{P}^2(r^2(LH)) = 0$, $\underline{P}^1(r^1(LH)) = \underline{P}^2(r^2(HL)) = 0.5$, and $\overline{P}^1(r^1(LH)) = \overline{P}^2(r^2(HL)) = 0.75$;

- $r^k(LL) = f_4^k(e^k) = -(e^k)^2$, this function represents the payoff each player k gets when they and the other player choose to contribute low effort. The lower and upper previsions of $r^k(LL)$ are $\underline{P}^k(r^k(LL)) = 0$ and $\overline{P}^k(r^k(LL)) = -0.25$.

The resulting vacuous previsions should be refined using an elicitation model and used in the enhanced GRR matrix. This matrix will return a convex set of correlated equilibria based on the uncertain outcomes' lower and upper utility previsions. Furthermore, in the repeated version of this game, dynamic updating can be used to adjust these utility previsions to the behaviour of each player. Dynamic updating will lead to an improved set of correlated equilibria by providing a more accurate utility assessment for each outcome. That is because Equations (4.9) and (4.10) adjust elicited previsions to newly observed utilities based on each player's delivered efforts.

7.2 Conclusions

Ignorance over a game's outcome has had limited treatment in game theory's literature, even though such cases are empirically evident. Here, we proposed the enhanced GRR matrix complemented with an elicitation model and NPUI-based dynamic updating as a normative solution to those situations.

We developed an enhanced version of Nau's revealed-rules matrix that allows

solving the correlated equilibria of a game that includes unknown outcomes. We conceptualised and built our solution using risk-neutral players, then extended it to cover the more generic case of risk-averse players.

We showed that the enhanced revealed-rules matrix reveals information that is not observable within the game's utility matrix. Specifically, it allows a player to include their lower and upper expectations towards the utility of an unknown outcome.

We discussed rationality requirements for lower and upper expectations and suggested an elicitation model allowing a player to assess them coherently. This elicitation model uses pre-existing beliefs as information, allowing cases of extreme ignorance, *i.e.* when a player has no information regarding an outcome.

We demonstrated that the enhanced matrix complemented with the elicitation model creates a complete framework for assessing a game's correlated equilibria under ignorance. An extensive example of a modified version of 'battle of the sexes', which includes ignorance over outcomes, was provided.

We examined repeated games and introduced a non-parametric updating model that allows a player to dynamically adjust an uncertain outcome's lower and upper expectations whenever it is experienced. To validate this updating model, we used a case of extreme ignorance and simulated a modified version of the game 'matching pennies'. The results showed that lower and upper previsions of the relevant outcome converge towards its linear prevision; that is, its actual expected payoff, which is initially unknown to the player.

We considered the suggested model's scalability and ability to support multiple sources of uncertainty. In such a case, the lower and upper previsions of unknown outcomes should not be limited to local assessments, *i.e.* independently

assign lower and upper previsions to each source of uncertainty. Instead, structural assessments should be considered, ensuring a coherent relation between lower (upper) previsions across all gambles.

We studied the impact of ignorance on specific game properties. Precisely, zero-sum and symmetry properties. We showed that uncertainty over outcomes could alter these properties and change the analysed structure of the relevant games. Furthermore, we leveraged the enhanced revealed-rules matrix to provide a solution that could explicitly force these properties.

Finally, we discussed the implication of ignorance on utility diagnostics and developed a new approach that allows assessing the impact of information patterns on strategic interactions. It allows computing the expected and actual value of a piece of information. We demonstrated this approach using the modified version of the game ‘battle of the sexes’.

In conclusion, the application of the proposed model could be explored in several fields of study. For example, in Artificial Intelligence, if a payoff matrix contains unobserved outcomes, it could be replaced with an enhanced revealed-rules matrix that includes players’ elicited lower and upper previsions towards these outcomes. Or within Economics, in rivalry or alliance situations where ignorance can prevail. For instance, when companies invest in a new market and unforeseen events arise, this could cause a conflict of interest. In this case, stakeholders could assess the value of further information through the suggested utility diagnostics approach and update their preferences using the dynamic updating mechanism.

List of Symbols

U	Utility function
ϕ	Outcome of a game-play
k	Player in a game
Λ	Player's set of all available alternatives in a game
a	Player's available alternative in a game
K	Total number of players in a game
m	Player's total number of available alternatives in a game
Φ	Set of all possible outcomes of a game
c	Risk aversion degree
r	Payoff function/vector
$\vec{\varepsilon}$	Player's mixed strategy
ε	Mixed strategy profile
E	Player's expected payoff function
$\vec{\rho}$	A correlated strategy
N	Total number of outcomes in a game

I	Set of all joint probability distributions that are independent between players
\underline{P}	Lower prevision
\overline{P}	Upper prevision
\mathcal{X}	Set of all possible outcomes of an experiment
f	Reward function referred to as a gamble
δ	Indicator of events function
\mathcal{A}	subset of possible outcomes of an experiment
e	An event
\mathbf{M}	Game revealed-rules matrix for risk neutral players
\dot{U}	Utility function's derivative
σ	Risk value
θ	Transaction's payoff
\mathbf{M}^*	Game revealed-rules matrix for risk averse players
n	Number of observations
u	Observed utility value
\hat{u}	Pre-observed utility value
F	Set of reward functions referred to as a gambles
$\underline{\mathcal{E}}$	Natural extension of the lower prevision of a given set of gambles
x	Outcome of an experiment
D	Set of almost-desirable gambles

\mathcal{E}	Natural extension of the linear prevision of a given set of gambles
w	Worst payoff of a gamble
b	Best payoff of a gamble
g	Reward function referred to as a gamble
\mathcal{Y}	Set of all possible outcomes of an experiment
y	Outcome of an experiment
Σ	Set of permutations
π	a permutation
Π	Random variable that represent unknown information statements
\mathcal{I}	Set of all possible information statements
ι	Information statement
V	Value of information function

Bibliography

- [1] K. J. Arrow. The theory of risk aversion. In *Essays in the Theory of Risk-bearing*, Markham economics series, chapter 3. North-Holland, 1971.
- [2] S. V. Astanin and N. K. Zhukovskaja. Using game theory and fuzzy logic to determine the dominant motivation cognitive agent. *American Association for Science and Technology*, 2:207–214, 2015.
- [3] S. Athey and K. Bagwell. Optimal collusion with private information. *RAND Journal of Economics*, 32:428–465, 2001.
- [4] T. Augustin and F. P. A. Coolen. Nonparametric predictive inference and interval probability. *Journal of Statistical Planning and Inference*, 124(2):251 – 272, 2004.
- [5] T. Augustin, F. P. A. Coolen, G. De Cooman, and M. C. M. Troffaes, editors. *Introduction to Imprecise Probabilities*. Wiley Series in Probability and Statistics. John Wiley and Sons, 2014.
- [6] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67 – 96, 1974.
- [7] R. J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- [8] R. J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55(1):1–18, 1987.
- [9] R.M. Baker, T. Coolen-Maturi, and F.P.A. Coolen. Nonparametric predictive inference for stock returns. *Journal of Applied Statistics*, 44(8):1333–1349, 2017.
- [10] D. Banks, V. Gallego, R. Naveiro, and D. Ríos Insua. Adversarial risk analysis: An overview. *WIREs Computational Statistics*, 14(1):e1530, 2022.

- [11] J. M. Bernardo. Expected information as expected utility. *The Annals of Statistics*, 7(3):686–690, 1979.
- [12] D. Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36, 1954.
- [13] D. Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, Berkeley, Calif., 1951. University of California Press.
- [14] E. Borel. *Traité du calcul des probabilités et de ses applications*. Number v. 4 in *Traité du calcul des probabilités et de ses applications*. Gautier-Villars et cie, 1931.
- [15] E. Borel and J. Ville. *Applications aux jeux de hasard*. *Traité de Calcul de Probabilités et de ses applications: Applications diverses et conclusion*. Gauthier-Vilars, 1953.
- [16] Z. Cao and X. Yang. Symmetric games revisited. *Mathematical Social Sciences*, 95(C):9–18, 2018.
- [17] A. Chakeri, A. N. Dariani, and C. Lucas. How can fuzzy logic determine game equilibriums better? In *2008 4th International IEEE Conference Intelligent Systems*, volume 1, Sep. 2008.
- [18] B. Chiandotto. *Bayesian and Non-Bayesian Approaches to Statistical Inference: A Personal View*, pages 3–13. Springer International Publishing, Cham, 2014.
- [19] M.L.D. Chiara, K. Doets, D. Mundici, and J. van Benthem. *Structures and Norms in Science: Volume Two of the Tenth International Congress of Logic, Methodology and Philosophy of Science, Florence, August 1995*. Synthese Library. Springer Netherlands, 2013.
- [20] B. de Finetti. La prévision: Ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 17:1–68, 1937.
- [21] B. de Finetti. *Theory of Probability: A Critical Introductory Treatment*, volume 1 and 2 of *Probability and Statistics Series*. Wiley, 1974, 1975.
- [22] M. H. DeGroot. Uncertainty, information and sequential experiments. *Annals of Mathematical Statistics*, 33:404–519, 1962.
- [23] M. H. DeGroot. Changes in utility as information. *Theory and Decision*, 17:287–303, 1984.

- [24] A. P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *The Annals of Mathematical Statistics*, 38(2):325–339, 1967.
- [25] T. Denoeux and M. Masson. *Belief Functions: Theory and Applications*. Springer, Heidelberg, 05 2012.
- [26] C. Elkan. The paradoxical success of fuzzy logic. *IEEE Expert*, 9(4):3–49, Aug 1994.
- [27] D. Ellsberg. Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669, 1961.
- [28] B. Fares. Game-Theory-Under-Uncertainty, 2021. <https://github.com/bfares/Game-Theory-Under-Uncertainty>.
- [29] T. Gajdos, J. M. Tallon, and J. C. Vergnaud. Decision making with imprecise probabilistic information. *Journal of Mathematical Economics*, 40(6):647–681, 2004.
- [30] D. Gale. *The Theory of Linear Economic Models*. McGraw-Hill, New York, 1960.
- [31] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
- [32] M. M. Gupta and G. K. Knopf. Fuzzy logic and uncertainty in image processing. *Proceedings of SPIE*, 1:210, 1993.
- [33] J. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii part i. the basic model. *Management Science*, 14(3):159–182, 1967.
- [34] B. M. Hill. Posterior distribution of percentiles: Bayes' theorem for sampling from a population. *Journal of the American Statistical Association*, 63(322):677–691, 1968.
- [35] B. M. Hill. De Finetti's theorem, induction, and $A_{(n)}$ or Bayesian nonparametric predictive inference (with discussion). *Bayesian statistics*, 3:211–241, 1988.
- [36] B. M. Hill. Parametric models for $A_{(n)}$: Splitting processes and mixtures. *Journal of the Royal Statistical Society. Series B (Methodological)*, 55(2):423–433, 1993.
- [37] B. Houlding and F. P. A. Coolen. Adaptive utility and trial aversion. *Journal of Statistical Planning and Inference*, 141, 2011.

- [38] B. Houlding and F. P. A. Coolen. Nonparametric predictive utility inference. *European Journal of Operational Research*, 221:222–230, 2012.
- [39] C. Jansen, H. Blocher, T. Augustin, and G. Schollmeyer. Information efficient learning of complexly structured preferences: Elicitation procedures and their application to decision making under uncertainty. *International Journal of Approximate Reasoning*, 144:69–91, 2022.
- [40] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Cambridge University Press, 1993.
- [41] C. Kirkwood. *Notes on attitude toward risk taking and the exponential utility function*. Department of Management, Arizona State University, 1997.
- [42] A. N. Kolmogorov. *Foundations of the Theory of Probability(Trans)*. Chelsea Publishing Company, 1950.
- [43] H. W. Kuhn. Extensive games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(10):570–576, 1950.
- [44] G. Laffond, J. Laslier, and M. Le Breton. A theorem on symmetric two-player zero-sum games. *Journal of Economic Theory*, 72(2):426–431, 1997.
- [45] D. Li, J. Nan, and M. Zhang. Interval programming models for matrix games with interval payoffs. *Optimization Methods & Software*, 27:1–16, 02 2012.
- [46] S. T. Liu and C. Kao. Matrix games with interval data. *Computers & Industrial Engineering*, 56(4):1697–1700, 2009.
- [47] R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. Publications of the Bureau of applied social research of the Columbia University. John Wiley & Sons, Incorporated, 1957.
- [48] R.P.S. Mahler. *Statistical Multisource-multitarget Information Fusion*. Artech House information warfare library. Artech House, 2007.
- [49] J. M. Mendel. Uncertainty, fuzzy logic, and signal processing. *Signal Processing*, 80(6):913 – 933, 2000.
- [50] J. Merrick and G. S. Parnell. A comparative analysis of pra and intelligent adversary methods for counterterrorism risk management. *Risk Analysis*, 31(9):1488–1510, 2011.
- [51] A. Motro. Imprecision and uncertainty in database systems. In Patrick Bosc and Janusz Kacprzyk, editors, *Fuzziness in Database Management Systems*, pages 3–22, Heidelberg, 1995. Physica-Verlag HD.

- [52] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- [53] J. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
- [54] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [55] R. Nau. Joint coherence in games of incomplete information. *Management Science*, 38(3):374–387, 1992.
- [56] R. Nau. Imprecise probabilities in non-cooperative games. *ISIPTA 2011 - Proceedings of the 7th International Symposium on Imprecise Probability: Theories and Applications*, 01 2011.
- [57] R. Nau, S. Canovas, and P. Hansen. On the geometry of nash equilibria and correlated equilibria. *International Journal of Game Theory*, 32:443–453, 02 2004.
- [58] R. Nau and K. McCardle. Coherent behavior in noncooperative games. *Journal of Economic Theory*, 50(2):424 – 444, 1990.
- [59] M. J. Osborne. Strategic and extensive games. 06 2006.
- [60] M. J. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2009.
- [61] M. J. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press, Cambridge, USA, 1994.
- [62] C. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multi-player games. *J. ACM*, 55, 07 2008.
- [63] V. Pareto. *Manuale di economia politica*. Piccola biblioteca scientifica. Societa Editrice, 1906.
- [64] V. Pareto and A. Bonnet. *Manuel d'économie politique*. Number v. 2 in Bibliothèque internationale d'économie politique. Librairie générale de droit et de jurisprudence, 1963.
- [65] S. D. Parsons, P. Gymtrasiewicz, and M. Wooldridge. *Game Theory and Decision Theory in Agent-Based Systems*. Multiagent Systems, Artificial Societies, and Simulated Organizations. Springer US, 2012.

- [66] B. Peleg, J. Rosenmüller, and P. Sudhölter. *The canonical extensive form of a game form: symmetries*, pages 367–387. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.
- [67] W. Poundstone. *Prisoner’s Dilemma*. Doubleday, 1992.
- [68] J. W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1/2):122–136, 1964.
- [69] J. Quiggin. The value of information and the value of awareness. *Theory and Decision*, 80(2):167–185, 2016.
- [70] R. Radner. Repeated Partnership Games with Imperfect Monitoring and No Discounting. *The Review of Economic Studies*, 53(1):43–57, 01 1986.
- [71] R. Radner, R. Myerson, and E. Maskin. An example of a repeated partnership game with discounting and with uniformly inefficient equilibria. *Review of Economic Studies*, 53:59–70, 1986.
- [72] H. Raiffa and R. Schlaifer. *Applied Statistical Decision Theory*. Harvard Business School Publications. Division of Research, Graduate School of Business Administration, Harvard University, 1961.
- [73] B. Ristic. Target classification with imprecise likelihoods: Mahler’s approach. *IEEE Transactions on Aerospace and Electronic Systems*, 47(2):1530–1534, 2011.
- [74] B. Ristic, C. Gilliam, M. Byrne, and A. Benavoli. A tutorial on uncertainty modeling for machine reasoning. *Information Fusion*, 55:30 – 44, 2020.
- [75] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–87, 02 1989.
- [76] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 2020.
- [77] P. Smets. *The transferable belief model for quantified belief representation*, pages 267–301. Springer, 1998.
- [78] P. Smets. *Belief Functions: The Disjunctive Rule of Combination and the Generalized Bayesian Theorem*, pages 633–664. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [79] P. Smets and R. Kennes. The transferable belief model. *Artificial Intelligence*, 66(2):191–234, 1994.

- [80] P. Sudhölter, J. Rosenmüller, and B. Peleg. The canonical extensive form of a game form: Part ii. representation. *Journal of Mathematical Economics*, 33(3):299–338, 2000.
- [81] M. C.M. Troffaes. Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning*, 45(1):17–29, 2007.
- [82] A. W. Tucker. The mathematics of tucker: A sampler. *The Two-Year College Mathematics Journal*, 14(3):228–232, 1983.
- [83] T. Verma and A. Kumar. *Fuzzy Solution Concepts for Non-cooperative Games: Interval, Fuzzy and Intuitionistic Fuzzy Payoffs*. Studies in Fuzziness and Soft Computing. Springer International Publishing, 2020.
- [84] J. Von Neumann. Zur Theorie der Gesellschaftsspiele. (German) [On the theory of games of strategy]. *j-MATH-ANN*, 100:295–320, 1928.
- [85] J. Von Neumann and O. Morgenstern. *Theory Of Games And Economic Behavior*. Princeton University Press, 1944.
- [86] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, 1991.
- [87] M. R. Walls. Measuring and utilizing corporate risk tolerance to improve investment decision making. *Engineering Economist*, 50(4):361 – 376, 2005.
- [88] P. M. Williams. Indeterminate probabilities. In M. Przelecki, K. Szaniawski, R. Wójcicki, and G. Malinowski, editors, *Formal Methods in the Methodology of Empirical Sciences: Proceedings of the Conference for Formal Methods in the Methodology of Empirical Sciences, Warsaw, June 17–21, 1974*, pages 229–246, Dordrecht, Netherlands, 1974.
- [89] P. M. Williams. Coherence, strict coherence and zero probabilities. In *DLMPS '75: Proceedings of the Fifth International Congress of Logic, Methodology and Philosophy of Science, volume VI*, pages 29–33, 01 1975.
- [90] Y. Yin, F.P.A. Coolen, and T. Coolen-Maturi. An imprecise statistical method for accelerated life testing using the power-weibull model. *Reliability Engineering & System Safety*, 167:158 – 167, 2017.
- [91] L. A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1(1):3–28, 1978.
- [92] L. A. Zadeh. The role of fuzzy logic in the management of uncertainty in expert systems. *Fuzzy Sets and Systems*, 11(1):199 – 227, 1983.

- [93] E. Zermelo. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In *Proceedings of the Fifth International Congress Mathematics*, pages 501–504, Cambridge, 1913. Cambridge University Press.